

The Memory Complexity of Depth-Robust and Dispersed Graphs

Joël Alwen
IST Austria

Jeremiah Blocki
Purdue University

Krzysztof Pietrzak
IST Austria

September 20, 2016

Abstract

There is growing interest in the security community in functions which are moderately hard to compute on a normal single-processor machine, but which cannot be computed at a significantly lower cost on dedicated hardware. Such functions are required for password-hashing to prevent brute-force attacks implemented on custom circuits and for ASIC resistant proofs-of-work for cryptocurrencies. Towards this goal memory-hard functions (MHF) have been suggested, motivated by the fact that – unlike computation – memory cost on custom hardware is not much cheaper than on general architectures. Of particular interest in the context of password hashing are data-independent MHFs (iMHFs) as they enjoy a natural resistance to certain side channel attacks which might otherwise leak information about users’ passwords.

An iMHF can be specified by fixing a directed acyclic graph (DAG) G_n on n nodes of constant indegree and a single sink representing the dependency graph between intermediary calls to an underlying hash function during the computation of the iMHF. The quality of the memory-hard function is captured by two parameters on the pebbling complexity of G_n :

- The parallel cumulative pebbling complexity $\Pi_{cc}^{\parallel}(G_n)$ must be as high as possible (to ensure that the amortized cost of computing the function on dedicated hardware is dominated by the cost of memory).
- The sequential space-time pebbling complexity $\Pi_{st}(G_n)$ should be as close as possible to $\Pi_{cc}^{\parallel}(G_n)$ (to ensure that using many cores in parallel and amortizing over many instances does not give much of an advantage).

In this paper we construct a family of DAGs with best possible parameters in an asymptotic sense, i.e., where $\Pi_{cc}^{\parallel}(G_n) = \Omega(n^2 / \log(n))$ (which matches a known upper bound) and $\Pi_{st}(G_n)$ is within a constant factor of $\Pi_{cc}^{\parallel}(G_n)$.

Our analysis relies on a new connection between depth-robustness (DR) of DAGs – a well studied combinatorial property – and their pebbling complexities. We show that high DR is *sufficient* for high Π_{cc}^{\parallel} . This contrasts the recent results of Alwen and Blocki (CRYPTO’16) showing that high DR is *necessary*. Together these results fully characterizes DAGs with high Π_{cc}^{\parallel} in terms of DR. Along the way we also give a new tool for reducing the indegree of a DAG without losing too much in Π_{cc}^{\parallel} which may be of interest beyond this work.

Complementing these results, in the second part of this work we provide new upper and lower bounds on the Π_{cc}^{\parallel} of several important iMHF candidates from the literature. We prove a $\tilde{\Omega}(n^{5/3})$ lower bound on Π_{cc}^{\parallel} for Argon2i, and a $\tilde{\Omega}(n^{1.5})$ lower bound for Catena and Balloon Hashing. The lower bound for Argon2i leverages our connection between Π_{cc}^{\parallel} and DR. To prove the lower bounds for Catena and Balloon Hashing we identify a graph property dubbed “dispersity”, and show how it also lower bounds Π_{cc}^{\parallel} .

Conversely we also describe a new class of pebbling attacks improving on those of Alwen and Blocki (CRYPTO’16). We apply them to Argon2i and Catena and upper bound their Π_{cc}^{\parallel} to be $O(n^{1.71})$ and $O(n^{1.62})$ respectively. As is the case with the Alwen-Blocki attacks we expect the new attack to be applicable to a wide variety of iMHFs from the literature and to have significantly lower memory complexity in practice than indicated by its asymptotic analysis.

1 Introduction

Moderately hard functions. Functions which are “moderately” hard to compute have found a variety of practical applications including password hashing, key-derivation and for proofs of work. In the context of password hashing, the goal is to minimize the damage done by a security breach where an adversary learns the password file; Instead of storing $(login, password)$ tuples in the clear, one picks a random salt and stores a tuple $(login, f(password, salt), salt)$, where $f(\cdot)$ is a moderately hard function $f(\cdot)$. This comes at a price, the server verifying a password must evaluate $f(\cdot)$, which thus cannot be too hard. On the other hand, if a tuple $(login, y, salt)$ is leaked, an adversary who tries to find the password by a dictionary attack must evaluate $f(\cdot)$ for every attempt. A popular moderately hard function is PBKDF2 (Password Based Key Derivation Function 2) [Kal00], which basically just iterates a cryptographic hash function H several times (1024 is a typical value).

Unfortunately a moderately hard function like PBKDF2 offers much less protection against adversaries who can build customized hardware to evaluate the underlying hash function than one would hope for. The reason is that the cost of computing a hash function H like SHA256 or MD5 on an ASIC (Application Specific Integrated Circuit) is orders of magnitude smaller than the cost of computing H on traditional hardware [DGN03, NB⁺15].

Memory-Bound and Memory-Hard Functions. [ABW03] recognized that cache-misses are more egalitarian than computation, in the sense that they cost about the same on different architectures. They propose “memory-bound” functions, which are functions that will incur many expensive cache-misses, this idea was further developed by [DGN03].

Along similar lines, Percival [Per09] observes that unlike computation, memory costs tend to be relatively stable across different architectures, and suggests to use memory-hard functions (MHF) for password hashing. [Per09] also introduced the `scrypt` MHF which has found a variety of applications in practice. Very recently it has been proven to indeed offer optimal time/space trade-offs in the random oracle model [ACP⁺16, ACK⁺16a].

MHFs come in two flavours, data-dependent MHFs (dMHF) such as `scrypt`, and data independent MHFs (iMHF). The former are potentially easier to construct and allow for more extreme memory-hardness [ACP⁺16, AB16a], but they leave open the possibility of side-channel attacks [FLW13], thus iMHFs are preferable when the inputs are sensitive, as in the case of password hashing. We shortly discuss the state of the art for dMHFs at the end of this section.

iMHF as Graphs. An iMHF comes with an algorithm that computes the function using a fixed memory access pattern. In particular the pattern is independent of the input. Such functions can thus be described by a directed acyclic graph (DAG) G , where each node v of the graph corresponds to some intermediate value ℓ_v that appears during the computation of the function, and the edges capture the computation: if ℓ_v is a function of previously computed values $\ell_{i_1}, \dots, \ell_{i_\delta}$, then the nodes i_1, \dots, i_δ are parents of v in G . For an iMHF F , we’ll denote with $G(F)$ the underlying graph. For example $G(\text{PBKDF2})$ is simply a path.

Graph Labeling Functions. Not only can an iMHF be captured by a graph as just outlined, we will actually construct iMHFs by first specifying a graph, and then defining a “labeling function” on top of it: Given a graph G with vertex set $V = [n]$, a hash function $H : \{0, 1\}^* \rightarrow \{0, 1\}^w$ and some input x define the labeling of the vertices of G as follows: a source (a vertex with indegree 0) has label $\ell_i = H(i, x)$, a vertex i with parents $i_1 < i_2 < \dots < i_\delta$ has label $\ell_i(x) = H(i, \ell_{i_1}(x), \dots, \ell_{i_\delta}(x))$. For a DAG G with a unique sink s we define the labeling function of G as $f_G(x) = \ell_s(x)$. Note that using the convention from the previous paragraph, we have $G(f_G) = G$.

The Black Pebbling Game One of the main techniques for analyzing iMHF is to use pebbling games played on graphs. First introduced by Hewitt and Paterson [HP70] and Cook [Coo73] the (sequential) black pebbling game (and its relatives) have been used to great effect in theoretical computer science. Some

early applications include space/time trade-offs for various computational tasks such as matrix multiplication [Tom78], the FFT [SS78, Tom78], integer multiplication [SS79b] and solving linear recursions [Cha73, SS79a]. More recently, pebbling games have been used for various cryptographic applications including proofs of space [DFKP15, RD16], proofs of work [DNW05, MMV13], leakage-resilient cryptography [DKW11a], garbled circuits [HJO⁺16], one-time computable functions [DKW11b], adaptive security proofs [HJO⁺16, JW16] and memory-hard functions [FLW13, AS15, AB16a, AGK⁺16]. It's also an active research topic in proof complexity (cf. The survey on <http://www.csc.kth.se/~jakobn/research/PebblingSurveyTMP.pdf>).

The black pebbling game is played over a fixed directed acyclic graph (DAG) $G = (V, E)$ in rounds. The goal of the game is to pebble all sink nodes of G (not necessarily simultaneously). Each round $i \geq 1$ is characterized by its pebbling configuration $P_i \subseteq V$ which denotes the set of currently pebbled nodes. Initially $P_0 = \emptyset$, i.e., all nodes are unpebbled. P_i is derived from the previous configuration P_{i-1} according to two simple rules. (1) A node v may be pebbled (added to P_i) if, in the previous configuration all of its parents were pebbled, i.e., $\text{parents}(v) \subseteq P_{i-1}$. (2) A pebble can always be removed from P_i . In the sequential version rule (1) may be applied at most once per round while in the parallel version no such restriction applies. A sequence of configurations $P = (P_0, P_1, \dots)$ is a (sequential) pebbling of G if it adheres to these rules and each sink node of G is contained in at least one configuration.

From a technical perspective, in this paper we investigate upper and lower bounds on various pebbling complexities of graphs, as they can be related to the cost of evaluating the “labeling function” f_G (to be defined below) in various computational models. In particular, let \mathcal{P}_G and \mathcal{P}_G^\parallel denote all valid sequential and parallel pebbblings of G , respectively. We are interested in the *parallel* cumulative pebbling complexity of G , denoted $\Pi_{cc}^\parallel(G)$, and the sequential space-time complexity of G , denoted $\Pi_{st}(G)$, which are defined as

$$\Pi_{cc}^\parallel(G) = \min_{(P_1, \dots, P_t) \in \mathcal{P}_G^\parallel} \sum_{i=1}^t |P_i| \quad \Pi_{st}(G) = \min_{(P_1, \dots, P_t) \in \mathcal{P}_G} t \cdot \max_i (|P_i|)$$

A main technical result of this paper is a family of graphs with high (in fact, as we'll discuss below, maximum possible) Π_{cc}^\parallel complexity, and where the Π_{st} complexity is not much higher than the Π_{cc}^\parallel complexity.¹ Throughout, we'll denote with \mathbb{G}_n the set of all DAGs on n vertices and with $\mathbb{G}_{n,d} \subseteq \mathbb{G}_n$ the DAGs where each vertex has indegree at most d .

Theorem 1.1 *There exists a family of DAGs $\{G_n \in \mathbb{G}_{n,2}\}_{n \in \mathbb{N}}$ where*

1. *parallel cumulative pebbling complexity*

$$\Pi_{cc}^\parallel(G_n) \in \Omega(n^2 / \log(n))$$

2. *and where the sequential space-time complexity matches the parallel cumulative pebbling complexity up to a constant*

$$\Pi_{st}(G_n) \in O(n^2 / \log(n))$$

The lower bound on Π_{cc}^\parallel in item 1. above is basically optimal due to the following bound from [AB16a].²

Theorem 1.2 ([AB16a, Thm. 8]) *For any constant $\epsilon > 0$ and sequence of DAGs $\{G_n \in \mathbb{G}_{n,\delta_n}\}_{n \in \mathbb{N}}$ it holds that*

$$\Pi_{cc}^\parallel(G_n) = o\left(\frac{\delta_n n^2}{\log^{1-\epsilon}}\right).$$

In particular if $\delta_n = O(\log^{1-\epsilon})$ then $\Pi_{cc}^\parallel(G_n) = o(n^2)$, and

$$\text{if } \delta_n = \Theta(1) \text{ then } \Pi_{cc}^\parallel(G_n) = o(n^2 / \log^{1-\epsilon}(n)) \tag{1}$$

¹Note that $\Pi_{cc}^\parallel(G) \leq \Pi_{st}(G)$ as parallelism can only help, and space-time complexity (i.e., number of rounds times the size of the largest state) is always higher than cumulative complexity (the sum of the sizes of all states).

²The statement below is obtained from the result in [AB16a] by treating the core-memory ratio as a constant and observing that, trivially, at most n pebbles are on G during a balloon phase and at most n pebbles are placed in one step during a balloon phase.

Pebbling vs. Memory-hardness. The reason to focus on the graph $G = G(F)$ underlying an iMHF F is that clean combinatorial properties of G – i.e., bounds on the pebbling complexities – imply both upper and lower bounds on the cost of evaluating F in various computational models. For upper bounds (i.e., attacks), no further assumption on F are required to make the transition from pebbling to computation cost. For lower bounds, we have to assume that there’s no “shortcut” in computing F , and the only way is to follow the evaluation sequence as given by G . Given the current state of complexity theory, where not even superlinear lower bounds on evaluating any function in \mathcal{NP} are known, we cannot hope to exclude such shortcuts unconditionally. Instead, we assume that the underlying hash function H is a random oracle and circuits are charged unit cost for queries to the random oracle.

For our lower bounds, we must insist on G having constant indegree. The reason is that in reality H must be instantiated with some cryptographic hash function like SHA1, and the indegree corresponds to the input length on which H is invoked. To evaluate H on long inputs, one would typically use some iterated construction like Merkle-Damgard, and the assumption that H behaves like a black-box that can only be queried once the entire input is known would be simply wrong in this case.

As advocated in [AS15], bounds on $\Pi_{cc}^{\parallel}(G)$ are a reasonable approximation for the cost of evaluating f_G in dedicated hardware, whereas a bound on $\Pi_{st}(G)$ gives an upper bound on the cost of evaluating f_G on a single processor machine. The reason [AS15] consider cumulative complexity for lower and space-time complexity for the upper bound is that when lower bounding the cost of evaluating f_G we do want to allow for amortization of the cost over arbitrary many instances,³ whereas for our upper bound we don’t want to make such an assumption. The reason we consider parallel complexity for the lower and only sequential for the upper bound is due to the fact that an adversary can put many (cheap) cores computing H on dedicated hardware, whereas for the upper bound we only want to consider a single processor machine.

If $\Pi_{cc}^{\parallel}(G)$ is sufficiently larger than $|V(G)|$ (in Theorem 1.1 it’s almost quadratic), then the cost of evaluating f_G in dedicated hardware is dominated by the memory cost. As memory costs about the same on dedicated hardware and general purpose processors, if our G additionally satisfies $\Pi_{cc}^{\parallel}(G) \approx \Pi_{st}(G)$, then we get a function f_G whose evaluation on dedicated hardware is not much cheaper than evaluating it on an off the shelf machine (like a single core x86 architecture). This is exactly what the family from Theorem 1.1 achieves. We elaborate on these computational models and how they are related to pebbling in Appendix A.

On the positive side, previous to this work, the construction with the best asymptotic bounds was due to [AS15] and achieved $\Pi_{cc}^{\parallel}(G_n) \in \Omega(n^2 / \log^{10}(n))$. However the exponent 10 (and the complexity of the construction) makes this an uninteresting for practical purposes.

On the negative side [AB16a, ACK⁺16a, AB16b] have broken many popular iMHFs in a rather strong asymptotic sense. For example, in [AB16a], the graph underlying Argon2i-A [BDK16], the winner of the recent password hashing competition⁴, was shown to have Π_{cc}^{\parallel} complexity $\tilde{O}(n^{1.75})$. For the recently proposed Balloon-Hashing [CGBS16] and for Catena [FLW13] the upper bound is $O(n^{1.67})$ is shown in [AB16a]. In [AB16b] these results were extended to show that Argon2i-B [BDKJ16] has $\Pi_{cc}^{\parallel}(G) = O(n^{1.8})$. Moreover [AB16b] show that for random instances of these functions (which is how they are supposed to be used in practice) the attacks actually have far lower Π_{cc}^{\parallel} than these asymptotic analyses indicate.

A New Generic Attack and Its Applications. In this work we improve on the attacks of [BK15, AS15, AB16a] (Section 6). We give a new parallel pebbling strategy for pebbling DAGs which lack a generalization of depth-robustness. Next we investigate this property for the case of Argon2i-A, Balloon-Hashing and Catena to obtain new upper bounds on their respective Π_{cc}^{\parallel} . For example we further improve the upper

³ Π_{cc}^{\parallel} satisfies a direct product property: pebbling k copies of G cost k times as much as pebbling G , i.e., $\Pi_{cc}^{\parallel}(G^k) = k \cdot \Pi_{cc}^{\parallel}(G)$, but this is not true for Π_{st} complexity.

⁴The Argon2 specification [BDK16] has undergone several revisions all of which are regularly referred to as “Argon2.” To avoid confusion we follow [AB16b] and use Argon2i-A [BDK15] to denote the version of Argon2i from the password hashing competition [PHC] and we use Argon2i-B [BDKJ16] to refer the version of Argon2 that is currently being considered for standardization by the Cryptography Form Research Group (CFRG) of the IRTF. We conjecture that the techniques introduced in this paper could also be used to establish tighter bounds for Argon2i-B in addition to Argon2i-A and Balloon-Hashing. However, we leave this as an open challenge for future work.

bound on Π_{cc}^{\parallel} for Argon2i-A from $\tilde{O}(n^{1.75})$ to $O(n^{1.708})$.

New Security Proofs. Complementing these results, in Section 5, we give the first security proofs for a variety of iMHFs. Hitherto the only MHF with a full security proof in a parallel computational model was [AS15] which employed relatively construction specific techniques. When restricted to sequential computation the results of [LT82, AS15] show that Catena has Π_{st} complexity $\Omega(n^2)$. Similar results are also shown for Argon2i-A and Balloon Hashing in [CGBS16].

In this work we introduce two new techniques for proving security of iMHFs. In the case of Argon2i-A we analyze its depth-robustness to show that its Π_{cc}^{\parallel} is at least $\tilde{\Omega}(n^{5/3})$. The second technique involves a new combinatorial property called dispersion which we show to imply lower bounds on the Π_{cc}^{\parallel} of a graph. We investigate the dispersion properties of the Catena and Balloon Hashing variants to show their Π_{cc}^{\parallel} to be $\tilde{\Omega}(n^{1.5})$. Previously no (non-trivial) lower bounds on Π_{cc}^{\parallel} were known for Argon2i-A, Catena or Balloon Hashing. Interestingly, our results show that Argon2i-A has better asymptotic security guarantees than Catena since $\Pi_{cc}^{\parallel} = \Omega(n^{5/3})$ for Argon2i-A and $\Pi_{cc}^{\parallel} = o(n^{1.619})$.

While these lower bounds are significantly worse than what we might ideally hope for in a secure iMHF (e.g., $\Pi_{cc}^{\parallel} \geq \Omega(n^2/\log(n))$), we observe that, in light of our new attacks in Section 6, they are nearly tight. Unfortunately, together with the bounds on the sequential complexity of these algorithms our results do highlight a large asymptotic gap between the memory needed when computing the functions on parallel vs. sequential computational devices.

Depth-Robust Graphs. The results in this work rely on a new connection between the depth-robustness of a DAG and its Π_{cc}^{\parallel} complexity. A DAG G is (e, d) -depth-robust if, after removing any subset of at most e nodes there remains a directed path of length at least d . First investigated by Erdős, Graham and Szemerédi [EGS75], several such graphs enjoying low indegree and increasingly extreme depth-robustness have been constructed in the past [EGS75, PR80, Sch82, Sch83, MMV13] mainly in the context of proving lower-bounds on circuit complexity and turing machine time. Depth-robustness has been used as a key tool in the construction of cryptographic objects like proofs of sequential work [MMV13] and memory-hard functions [AS15].

Depth-Robustness and Π_{cc}^{\parallel} . While the flavour of the results in this work are related to those of [AS15] the techniques are rather different. As mentioned above already, they stem from a new tight connection between depth-robustness and Π_{cc}^{\parallel} . A special case of this connection shows that if G is (e, d) -depth-robust, then its Π_{cc}^{\parallel} can be lower bounded as

$$\Pi_{cc}^{\parallel}(G) \geq e \cdot d .$$

This complements a result from [AB16a], who give a pebbling strategy which is efficient for graphs of low depth-robustness (we give the exact statement in Theorem 2.5 below), thus a DAG has high Π_{cc}^{\parallel} if and only if it is very depth-robust.

Moreover, we give a new tool for reducing the indegree of a DAG while not reducing the Π_{cc}^{\parallel} of the resulting graph (in terms of its size). Together these results directly have some interesting consequences

- The family of DAGs $\{G_n \in \mathbb{G}_{n, \log(n)}\}_{n \in \mathbb{N}}$ from Erdős et al. [EGS75] have optimally high $\Pi_{cc}^{\parallel}(G_n) \in \Omega(n^2)$.
- Using our indegree reduction we can turn the above family of $\log(n)$ indegree into a family of indegree 2 DAGs $\{G'_n \in \mathbb{G}_{(n, 2)}\}_{n \in \mathbb{N}}$ with $\Pi_{cc}^{\parallel}(G'_n) \in \Omega(n^2/\log(n))$, which by Theorem 1.2 is optimal for constant indegree graphs.

Data-Dependent MHFs. One can naturally extend the Π_{cc}^{\parallel} notion also to “dynamic” graphs – where some edges are only revealed as some nodes are pebbled – in order to analyze data-dependent MHFs (dMHF) like `script`. In this model, [ACK⁺16b] show that $\Pi_{cc}^{\parallel}(\text{script}) = \Omega(n^2/\log^2(n))$. Unfortunately unlike for iMHFs, for dMHFs we do not have a proof that a lower bound on Π_{cc}^{\parallel} implies roughly the same lower bound on the cumulative memory complexity in the random oracle model.⁵ Recently a “direct” proof (i.e., avoiding pebbling arguments) – showing that `script` has optimal cumulative memory complexity $\Omega(n^2)$ – has been announced, note that this bound is better than what we can hope to achieve for iMHFs (as stated in Theorem 1.2). Unfortunately, the techniques that have now been developed to analyse dMHFs seem not to be useful for the iMHF setting.

2 Pebbling Complexities and Depth-Robustness of Graphs

We begin by fixing some common notation. We use the sets $\mathbb{N} = \{0, 1, 2, \dots\}$, $\mathbb{N}^+ = \{1, 2, \dots\}$, and $\mathbb{N}_{\geq c} = \{c, c+1, c+2, \dots\}$ for $c \in \mathbb{N}$. Further, we also use the sets $[c] := \{1, 2, \dots, c\}$ and $[b, c] = \{b, b+1, \dots, c\}$ where $b \in \mathbb{N}$ with $b \leq c$. For a set of sets $A = \{B_1, B_2, \dots, B_z\}$ we use the notation $\|A\| := \sum_i |B_i|$.

2.1 Depth-Robust Graphs

We say that a directed acyclic graph (DAG) $G = (V, E)$ has *size* n if $|V| = n$. A node $v \in V$ has indegree $\delta = \text{indeg}(v)$ if there exist δ incoming edges $\delta = |(V \times \{v\}) \cap E|$. More generally, we say that G has indegree $\delta = \text{indeg}(G)$ if the maximum indegree of any node of G is δ . A node with indegree 0 is called a source node and one with no outgoing edges is called a sink. We use $\text{parents}_G(v) = \{u \in V : (u, v) \in E\}$ to denote the parents of a node $v \in V$. In general, we use $\text{ancestors}_G(v) = \bigcup_{i \geq 1} \text{parents}_G^i(v)$ to denote the set of all ancestors of v — here, $\text{parents}_G^2(v) = \text{parents}_G(\text{parents}_G(v))$ denotes the grandparents of v and $\text{parents}^{i+1}(v) = \text{parents}_G(\text{parents}^i(G))$. When G is clear from context we will simply write parents (ancestors). We denote the set of all sinks of G with $\text{sinks}(G) = \{v \in V : \nexists (v, u) \in E\}$ — note that $\text{ancestors}(\text{sinks}(G)) = V$. We often consider the set of all DAGs of equal size $\mathbb{G}_n = \{G = (V, E) : |V| = n\}$ and often will bound the maximum indegree $\mathbb{G}_{n, \delta} = \{G \in \mathbb{G}_n : \text{indeg}(G) \leq \delta\}$. For directed path $p = (v_1, v_2, \dots, v_z)$ in G its length is the number of nodes it traverses $\text{length}(p) := z$. The depth $d = \text{depth}(G)$ of DAG G is the length of the longest directed path in G .

We will often consider graphs obtained from other graphs by removing subsets of nodes. Therefor if $S \subset V$ then we denote by $G - S$ the DAG obtained from G by removing nodes S and incident edges. The following is a central definition to our work.

Definition 2.1 (Depth-Robustness) For $n \in \mathbb{N}$ and $e, d \in [n]$ a DAG $G = (V, E)$ is (e, d) -depth-robust if

$$\forall S \subset V \quad |S| \leq e \Rightarrow \text{depth}(G - S) \geq d.$$

We will make use of the following lemma due to Erdős, Graham and Szemerédi [EGS75], who showed how to construct a family of log indegree DAGs with extreme depth-robustness.

Theorem 2.2 ([EGS75]) For some fixed constants $c_1, c_2, c_3 > 0$ there exists an infinite family of DAGs $\{G_n \in \mathbb{G}_{n, c_3 \log(n)}\}_{n=1}^{\infty}$ such that G_n is $(c_1 n, c_2 n)$ -depth-robust.

2.2 Graph Pebbling

We fix our notation for the parallel graph pebbling game following [AS15].

⁵[ACK⁺16b] introduces a combinatorial conjecture, which if true, means that lower bounds on Π_{cc}^{\parallel} translate to cumulative memory complexity. At this point a strong variant of the conjecture has already been refuted, the state of the conjecture is updated in the eprint version [ACK⁺16c] of the paper.

Definition 2.3 (Parallel/Sequential Graph Pebbling) Let $G = (V, E)$ be a DAG and let $T \subseteq V$ be a target set of nodes to be pebbled. A pebbling configuration (of G) is a subset $P_i \subseteq V$. A legal parallel pebbling of T is a sequence $P = (P_0, \dots, P_t)$ of pebbling configurations of G where $P_0 = \emptyset$ and which satisfies conditions 1 & 2 below. A sequential pebbling additionally must satisfy condition 3.

1. At some step every target node is pebbled (though not necessarily simultaneously).

$$\forall x \in T \exists z \leq t : x \in P_z.$$

2. Pebbles are added only when their predecessors already have a pebble at the end of the previous step.

$$\forall i \in [t] : x \in (P_i \setminus P_{i-1}) \Rightarrow \text{parents}(x) \subseteq P_{i-1}.$$

3. At most one pebble placed per step.

$$\forall i \in [t] : |P_i \setminus P_{i-1}| \leq 1.$$

We denote with $\mathcal{P}_{G,T}$ and $\mathcal{P}_{G,T}^{\parallel}$ the set of all legal sequential and parallel pebblings of G with target set T , respectively. Note that $\mathcal{P}_{G,T} \subseteq \mathcal{P}_{G,T}^{\parallel}$. We will be mostly interested in the case where $T = \text{sinks}(G)$ and then will simply write \mathcal{P}_G and $\mathcal{P}_G^{\parallel}$.

Definition 2.4 (Time/Space/Cumulative Pebbling Complexity) The space, time, space-time and cumulative complexity of a pebbling $P = \{P_0, \dots, P_t\} \in \mathcal{P}_G^{\parallel}$ are defined as

$$\Pi_t(P) = t, \quad \Pi_s(P) = \max_{i \in [t]} |P_i|, \quad \Pi_{st}(P) = \Pi_t(P) \cdot \Pi_s(P) \quad \text{and} \quad \Pi_{cc}(P) = \sum_{i \in [t]} |P_i|.$$

For $\alpha \in \{s, t, st, cc\}$ and a target set $T \subseteq V$, the sequential and parallel pebbling complexities of G are defined as

$$\Pi_{\alpha}(G, T) = \min_{P \in \mathcal{P}_{G,T}} \Pi_{\alpha}(P) \quad \text{and} \quad \Pi_{\alpha}^{\parallel}(G, T) = \min_{P \in \mathcal{P}_{G,T}^{\parallel}} \Pi_{\alpha}(P).$$

When $T = \text{sinks}(G)$ we simplify notation and write $\Pi_{\alpha}(G)$ and $\Pi_{\alpha}^{\parallel}(G)$.

It follows from the definition that for $\alpha \in \{s, t, st, cc\}$ and any G the parallel pebbling complexity is always at most as high as the sequential, i.e., $\Pi_{\alpha}(G) \geq \Pi_{\alpha}^{\parallel}(G)$, and cumulative complexity is at most as high as space-time complexity, i.e., $\Pi_{st}(G) \geq \Pi_{cc}(G)$ and $\Pi_{st}^{\parallel}(G) \geq \Pi_{cc}^{\parallel}(G)$.

In this work we will consider constant in-degree DAGs $\{G_n \in \mathbb{G}_{n, \Theta(1)}\}_{n \in \mathbb{N}}$, and will be interested in the complexities $\Pi_{st}(G_n)$ and $\Pi_{cc}^{\parallel}(G_n)$ as these will capture the cost of evaluating the labelling function derived from G_n on a single processor machine (e.g. a x86 processor on password server) and amortized AT complexity (which is a good measure for the cost of evaluating the function on dedicated hardware), respectively.

Before we state our main theorem let us observe some simple facts. Every n -vertex graph can be pebbled in n steps, and we cannot have more than n pebbles on a n vertex graph, thus

$$\forall G_n \in \mathbb{G}_n : \Pi_{cc}^{\parallel}(G_n) \leq \Pi_{st}(G_n) \leq n^2$$

This upper bound is basically matched for the complete graph $K_n = (V = [n], E = \{(i, j) : 1 \leq i < j \leq n\})$

$$n(n-1)/2 \leq \Pi_{cc}^{\parallel}(K_n) \leq \Pi_{st}(K_n) \leq n^2$$

K_n has the desirable properties that its Π_{st} is within a constant factor to its Π_{cc}^{\parallel} complexity and moreover its Π_{cc}^{\parallel} complexity is maximally high. Unfortunately, K_n has very high indegree, which makes it useless for

our purpose to construct memory-hard functions. The path $Q_n = (V = [n], E = \{(i, i + 1) : 1 \leq i \leq n - 1\})$ on the other hand has indegree 1 and its Π_{st} is even exactly as large as its Π_{cc}^{\parallel} complexity. Unfortunately it has very low pebbling complexity

$$\Pi_{cc}^{\parallel}(Q_n) = \Pi_{st}(Q_n) = n$$

which means that in the labelling function we get from Q_n (which is basically PBKDF2 discussed in the introduction) the evaluation cost will not be dominated by the memory cost even for large n . As stated in Theorem 1.1, in this paper we construct a family of graphs $\{G_n \in \mathbb{G}_{n,2}\}_{n \in \mathbb{N}}$ which satisfies all three properties at once: (1) the graphs have indegree 2 (2) the parallel cumulative pebbling complexity is $\Pi_{cc}^{\parallel}(G_n) \in \Omega(n^2/\log(n))$, which by Theorem 1.2 is optimal for constant indegree graphs, and (3) $\Pi_{st}(G_n)$ is within a constant factor of $\Pi_{cc}^{\parallel}(G_n)$.

We will use the following result from [AB16a] which shows that DAGs that are not depth-robust also have low Π_{cc}^{\parallel} .

Theorem 2.5 ([AB16a, Thm. 2]) *Let $G_n \in \mathbb{G}_{n,\delta}$ such that G_n is not (e, d) -depth-robust. Then*

$$\Pi_{cc}^{\parallel}(G_n) = O\left(\min_{g \in [d, n]} \left\{ n \left(\frac{dn}{g} + \delta g + e \right) \right\}\right)$$

setting $g = \sqrt{\frac{dn}{\delta}}$ this simplifies to

$$\Pi_{cc}^{\parallel}(G) = O\left(n(\sqrt{dn\delta} + e)\right)$$

3 Depth-Robustness Implies High Π_{cc}^{\parallel}

In this section we state and prove a theorem which lowerbounds the Π_{cc}^{\parallel} of a given DAG G in terms of its depth robustness. An special interesting case of the general theorem is the following corollary

Corollary 3.1 (of Theorem 3.2) *Let G be an (e, d) -depth-robust DAG, then $\Pi_{cc}^{\parallel}(G) > ed$.*

We will give a proof of this corollary because it's simpler than for the general case but already contains the main ideas of the proof. See Appendix C for the proof of Theorem 3.2.

Proof of Corollary 3.1. Let (P_1, \dots, P_m) be a parallel pebbling of minimum complexity, i.e., $\sum_{i=1}^m |P_i| = \Pi_{cc}^{\parallel}(G)$. For any d , we'll show that there exists a set B of size $|B| \leq \Pi_{cc}^{\parallel}(G)/d$ such that there's no path of length d in $G - B$, or equivalently, G is not $(\Pi_{cc}^{\parallel}(G)/d, d)$ -depth-robust, note that this implies the corollary.

For $i \in [d]$ define $B_i = P_i \cup P_{i+d} \cup P_{i+2d} \dots$. We observe that by construction $\sum_{i=0}^{d-1} |B_i| = \Pi_{cc}^{\parallel}(G)$, so the size of the B_i 's is $\Pi_{cc}^{\parallel}(G)/d$ on average, and the smallest B_i has size at most this. Let B be the smallest B_i , as just outlined $|B| \leq \Pi_{cc}^{\parallel}(G)/d$.

It remains to show that $G - B$ has no path of length d . For this consider any path v_1, \dots, v_d of length d in G . Let j be minimal such that $v_d \in P_j$ (so v_d is pebbled for the first time in round j of the pebbling). It then must be the case that $v_{d-1} \in P_{j-1}$ (as to pebble v_d is round j there must have been a pebble on v_{d-1} in round $j - 1$). In round $j - 2$ either the pebble on v_{d-1} was already there, or there was a pebble on v_{d-2} . This argument shows that each of the pebbling configurations $\{P_{j-d+1}, \dots, P_j\}$ must contain at least one node from v_1, \dots, v_d . As B contains each d th Pebbling configuration, it has nonempty intersection with $\{P_{j-d+1}, \dots, P_j\}$, thus the path v_1, \dots, v_d is not contained entirely in $G - B$. \square

Our general Theorem 3.2 states that it remains expensive to pebble any large enough set of remaining nodes in a depth-robust graph even if we are permitted to first remove an arbitrary node set of limited size. Recall that $\Pi_{cc}^{\parallel}(G)$ is the parallel pebbling of minimal cumulative cost when pebbling all sinks of G , this requires pebbling all nodes of G at least once. Thus Corollary 3.1 follows from Theorem 3.2 below by setting $S = \emptyset$ letting $T = \text{sinks}(G)$ in which case $\sigma = 0$ and $\tau = 1$.

One application of this general theorem involves analyzing the cost of pebbling stacks of depth-robust graphs. For example if there are not enough pebbles on the graph at some point in time then there must be some layers with few pebbles. If we can then show that many of the nodes on those layers will eventually need to be (re)pebbled then we can use this lemma to show that the remaining pebbling cost incurred by these layers is large.

Theorem 3.2 *Let DAG $G = (V, E)$ be (e, d) -depth-robust and let $S, T \subset V$ such that*

$$|S| \leq \sigma n \quad , \quad |\text{ancestors}_{G-S}(T)| \geq \tau n \quad \text{and} \quad T \cap S = \emptyset$$

Then the cost of pebbling $G - S$ with target set T is $\Pi_{cc}^{\parallel}(G - S, T) > (e - \sigma n)(d - (1 - \tau)n)$.

Another immediate implication of Theorem 3.2 and Theorem 2.2 is that there is an infinite family of DAGs with maximal $\Pi_{cc}^{\parallel}(G) = \Omega(n^2)$ whose indegree scales with $\log n$. Note that this means that allowing indegree as small as $O(\log(n))$ is sufficient to get DAGs whose Π_{cc}^{\parallel} is within a constant factor of the n^2 upper bound on Π_{cc}^{\parallel} for any n vertex DAG.

Corollary 3.3 (of Theorem 3.2 and Theorem 2.2) *For some constants $c_1, c_2 > 0$ there exists an infinite family of DAGs $\{G_{n,\delta} \in \mathbb{G}_{n,\delta}\}_{n=1}^{\infty}$ with $\delta \leq c_1 \log(n)$ and $\Pi_{cc}^{\parallel}(G) \geq c_2 n^2$.*

This is optimal in the sense that for any family $\{\delta_n \in [n]\}_{n=1}^{\infty}$ and $\{J_n \in \mathbb{G}_{n,\delta_n}\}_{n=1}^{\infty}$ it holds that $\Pi_{cc}^{\parallel}(J_n) \in O(n^2)$. Moreover if $\delta_n = o(\log(n)/\log \log(n))$ then $\Pi_{cc}^{\parallel}(J_n) = o(n^2) = o(\Pi_{cc}^{\parallel}(G_n))$.

4 Indegree Reduction: Constant Indegree with Maximal Π_{cc}^{\parallel}

In this section we use the result from the previous section to show a new, more efficient, degree-reduction lemma. We remark that Lemma 4.1 is similar to [AS15, Lemma 9] in that both reductions replace high indegree nodes v in G with a path. However, we stress two key differences between the two results. First, our focus is on reducing the indegree while preserving depth-robustness. By contrast, [AS15, Lemma 9] focuses directly on preserving Π_{cc}^{\parallel} . Second, we note that the guarantee of [AS15, Lemma 9] is weaker in that it yields a reduced indegree graph G' whose size grows by a factor of indeg ($n' \leq n \times \text{indeg}$) while Π_{cc}^{\parallel} can drop by a factor of indeg — [AS15, Lemma 9] shows that $\Pi_{cc}^{\parallel}(G') \geq \frac{\Pi_{cc}^{\parallel}(G)}{\text{indeg}-1}$. By contrast, setting $\gamma = \text{indeg}$ in Lemma 4.1 yields a reduced indegree graph G' whose size grows by a factor of $2 \times \text{indeg}$ ($n' \leq 2n \times \text{indeg}$) and better depth-robustness $(e', d') = (e, d \times \text{indeg})$. In particular, when we apply Corollary 3.1 the lower-bound $\Pi_{cc}^{\parallel}(G') \geq ed \times \text{indeg}$ improves by a factor of indeg when compared with the original graph G .

Lemma 4.1 *Let DAG G be (e, d) -depth-robust. For $\gamma \in \mathbb{Z}_{\geq 0}$ there exists $(e, d\gamma)$ -depth-robust DAG G' with*

$$\text{size}(G') \leq (\text{indeg}(G) + \gamma) \cdot \text{size}(G) \quad , \quad \text{indeg}(G') = 2 \quad \text{and} \quad \Pi_{st}(G') \leq \frac{\text{size}(G')^2}{\gamma}.$$

PROOF. Fix a $\gamma \in [n]$. We identify each node in V' with an element of the set $V \times [\delta + \gamma]$ and we write $\langle v, j \rangle \in V'$. For every node $v \in V$ with $\alpha_v := \text{indeg}(v) \in [0, \delta]$ we add the path $p_v = (\langle v, 1 \rangle, \langle v, 2 \rangle, \dots, \langle v, \alpha_v + \gamma \rangle)$ of length $\alpha_v + \gamma$. We call v the *genesis node* and p_v its *metanode*. In particular $V' = \cup_{v \in V} p_v$. Thus G has size at most $(\delta + \gamma)n$.

Next we add the remaining edges. Intuitively, for the i^{th} incoming edge (u, v) of v we add an edge to G' connecting the end of the metanode of u to the i^{th} node in the metanode of v . More precisely, for every $v \in V$, $i \in [\text{indeg}(v)]$ and edge $(u_i, v) \in E$ we add edge $(\langle u_i, \text{indeg}(u_i) + \gamma \rangle, \langle v, i \rangle)$ to E' . It follows immediately that G' has indegree (at most) 2.

Fix any node set $S \subset V'$ of size $|S| \leq e$. Then at most e metanodes can share a node with S . For each such metanode remove its genesis node in G . As G is (e, d) -depth-robust we are still left with a path p of length (at least) d in G . But that means that after removing S from G' there must remain a path p' in G' running through all the metanodes of p and $|p'| \geq |p|\gamma \geq d\gamma$. In other words G' is $(e, d\gamma)$ -depth-robust.

To see that $\Pi_{st}(G') \leq \text{size}(G')^2/\gamma$ we simply pebble G' in topological order. We note that we never need to keep more than one pebble on any metanode $p_v = (\langle v, 1 \rangle, \langle v, 2 \rangle, \dots, \langle v, \alpha_v + \gamma \rangle)$ with $\alpha_v = \text{indeg}(v)$. Once we pebble the last node $\langle v, \alpha_v + \gamma \rangle$ we can permanently discard any pebbles on the rest of p_v since $\langle v, \alpha_v + \gamma \rangle$ is the only node with outgoing edges. \square

Proof of Theorem 1.1. Theorem 1.1 follows by applying Lemma 4.1 to the family from Theorem 2.2 with $\gamma = \text{indeg} = \theta(\log n)$. We get that for some fixed constants $c_1, c_2 > 0$ there exists an infinite family of indegree 2 DAGs $\{G_n \in \mathbb{G}_{n,2}\}_{n=1}^\infty$ where G_n is $(c_1 n / \log n, c_2 n)$ -depth robust and $\Pi_{st}(G_n) \leq O(n^2 / \log(n))$. By Corollary 3.1 then $\Pi_{cc}^{\parallel}(G_n) > (c_1 c_2) n^2 / \log(n)$, which is basically optimal for constant indegree DAGs by Theorem 1.2. \square

5 Security Proofs of Candidate iMHFs

On the surface, in this and the next section we give both security proofs and nearly optimal attacks for several of the most prominent iMHF proposals. That is we show both lower and (relatively tight) upperbounds on their asymptotic memory-hardness in the PROM. However, more conceptually, we also introduce two new proof techniques for analyzing the depth-robustness of DAGs as well as a new very memory-efficient class of algorithms for pebbling a DAG improving on the techniques used in [AB16a]. Indeed for all candidates considered the attack in the next section is almost optimal in light of the accompanying security proofs in this section.

More specifically, in the first subsection we prove bounds for a class of random graphs which generalize the Argon2i-A construction [BDK16] and the Single-Buffer (SB) construction of [CGBS16]. To prove the lowerbound we use a simple and clean new technique for bounding the depth-robustness of a random DAG. Combined with Corollary 3.3 we could immediately obtain a lower bound of $\tilde{\Omega}(n^{1.5})$. We can improve the lower bound to $\tilde{O}(n^{5/3})$ by showing that we need to either (1) keep $\tilde{O}(n^{2/3})$ pebbles on the first half of the DAG during most pebbling steps, or (2) frequently repebble the first half of the DAG (which also has high cost).

In the second subsection we prove bounds for a family of layered graphs which generalize both of the Catena constructions [FLW13]. The upperbound for Catena is implied by the same result for the random DAGs. However the lowerbound is a direct proof which uses a new type of pebbling argument without going through the notion of depth-robustness.

5.1 Lowerbounding the CC of Random DAGs.

We begin by defining a (n, δ, w) -random DAG, the underlying DAGs upon which Argon2i-A and SB are based. The memory window parameter w specifies the intended memory usage and throughput of the iMHF — the cost of the naïve pebbling algorithm is $\Pi_{cc}^{\parallel}(\mathcal{N}) = wn$. In particular, a t -pass Argon2i-A iMHF is based on a $(n, 2, n/t)$ -random DAG. Similarly, a t -pass Single-Buffer (SB) iMHF [CGBS16] is based on a $(n, 20, n/t)$ -random DAG. In this section we focus on the $t = 1$ -pass variants of the Argon2i-A and [CGBS16] iMHFs.

Definition 5.1 ((n, δ, w)-random DAG) *Let $n \in \mathbb{N}$, $1 < \delta < n$, and $1 \leq w \leq n$ such that w divides n . An (n, δ, w) -random DAG is a randomly generated directed acyclic (multi)graph with n nodes v_1, \dots, v_n and with maximum in-degree δ for each vertex. The graph has directed edges (v_i, v_{i+1}) for $1 \leq i < n$ and random forward edges $(v_{r(i,1)}, v_i), \dots, (v_{r(i,\delta-1)}, v_i)$ for each vertex v_i . Here, $r(i, j)$ is independently chosen uniformly at random from the set $[\max\{0, i - w\}, i - 1]$.*

Security Lower Bound. To prove the lower bound we rely on a slightly stricter notion of depth robustness. Given a node v let $N(v, b) = \{v - b + 1, \dots, v\}$ denote a segment of b consecutive nodes ending at v and given a set $S \subseteq V(G)$ let $N(S, b) = \bigcup_{v \in S} N(v, b)$. We say that a DAG G is (e, d, b) -block depth-robust if for every set $S \subseteq V(G)$ of size $|S| \leq e$ we have $\text{depth}(G - N(S, b)) \geq d$. Lemma 5.3 shows that for any $e \geq \sqrt{n}$,

with high probability, a $(n, 2, n)$ -random DAG G will be (e, d, b) -block depth-robust with $d = \frac{n^2}{e^{2 \text{polylog}(n)}}$ and $b = n/(10e)$ (in contrast Lemma 6.3 states that G will be (e, d) -reducible with $d = \tilde{O}(n^2/e^2)$). Taking $e = \sqrt{n}$ in Corollary 3.1 immediately implies that $\Pi_{cc}^{\parallel}(G) = \tilde{\Omega}(n^{1.5})$. We can improve this bound with a slightly more sophisticated argument. In particular, we consider the cost incurred while pebbling a set $B \subseteq \{n/2 + 1, \dots, n\}$ of $e * \text{polylog}(n)$ consecutive nodes in the last half of G . Case 1: We keep $\Omega(e)$ pebbles on G (nodes $\{1, \dots, n/2\}$) while pebbling B_{first} , the first $|B|/2$ nodes of B (Total Cost: $\Omega(e^2) = \tilde{\Omega}(n^{4/3})$). Case 2: At some point in time t we have $|P_t| = o(e)$ pebbles on G while pebbling B_{first} . In this case, we will need to essentially repebble the entire first half of G before we can finish pebbling B_{last} , the last $|B|/2$ nodes of B . In particular, we can prove that (whp) every segment of $b = \tilde{\Omega}(n^{1/3})$ consecutive nodes in the first half of G has a directed edge to B_{last} . Now a similar argument to the proof of Theorem 3.2 shows that it will cost $\Omega(ed) = \tilde{\Omega}(n^{4/3})$ to repebble the necessary nodes in the first half of G — note that $G - \{n/2 + 1, \dots, n\}$ is a $(n/2, \delta, n/2)$ -random DAG and is thus block depth robust (whp) by Lemma 6.3. There are $n/(e * \text{polylog}(n)) = \tilde{\Omega}(n^{1/3})$ disjoint segments of $e * \text{polylog}(n)$ consecutive nodes in $\{n/2 + 1, \dots, n\}$. Thus, the total cost is $\tilde{\Omega}(n^{5/3})$.

Theorem 5.2 *Let G be a (n, δ, n) -random DAG then, except with probability $o(n^{-3})$, we have*

$$\Pi_{cc}^{\parallel}(G) = \tilde{\Omega}(n^{5/3}) .$$

We make a couple of observations.

1. The lower bound from Theorem 5.2 might be tight. Alwen and Blocki [AB16a] gave an attack \mathcal{A} such that $\Pi_{cc}^{\parallel}(\mathcal{A}) = O(n^{1.75} \delta \log n)$ for a (n, δ, t) -random DAG. We reduce the gap of $\tilde{O}(n^{1/12})$ by developing an improved recursive version of the attack of Alwen and Blocki [AB16a]. In particular, we show that for any $\epsilon > 0$ we have $\Pi_{cc}^{\parallel}(\mathcal{A}) = o(n^{1 + \sqrt{1/2 + \epsilon}}) = o(n^{1.708})$. Our modified attack also improves the upper bound for other iMHF candidates like Catena [FLW13].
2. Corollary 3.1 alone will not yield any meaningful lower bounds on the Π_{cc}^{\parallel} of the Catena iMHFs [FLW13]. In particular, the results from Alwen and Blocki [AB16a] imply that for any t -pass variant of Catena corresponding DAG is not (e, d) -depth robust for $ed \geq nt$ (typically, $t = O(\text{polylog}(n))$). However, we use an alternative techniques below to prove that $\Pi_{cc}^{\parallel}(G) = \Omega(n^{1.5})$ for the both Catena iMHFs and the Linear and Double-Buffer iMHFs of [CGBS16].

The proof of Theorem 5.2 relies on the next lemma. The proofs of Lemma 5.3 and Theorem 5.2 are in Appendix C.

Lemma 5.3 *For any $e \geq \sqrt{n}$ any any $\delta \geq 2$ a (n, δ, n) -random DAG will be $(e, \Omega(\frac{n^2}{e^2 \log(n)}), \frac{n}{100e})$ -block depth robust with probability $1 - o(n^{-3})$.*

5.2 Lowerbounding Dispersed Graphs

In this section we prove a lowerbound on the CC of a class of DAGs called dispersed graphs (defined bellow). Next we show that several of the iMHF constructions from the literature are based on such graphs. Thus we obtain proofs of security for each of these constructions (albeit for limited levels of security). In the subsequent section we give an upperbound on the CC of these constructions showing that the lowerbounds in this section are relatively tight.

Intuitively a (g, k) -dispersed DAG is a DAG ending with a path ϕ of length k which has widely dispersed dependencies. The following definitions make this concept precises.

Definition 5.4 (Dependencies) *Let $G = (V, E)$ be a DAG and $L \subseteq V$. We say that L has a (z, g) -dependency if there exist node disjoint paths p_1, \dots, p_z each ending in L and with length (at least) g .*

We are interested in graphs with long paths with many sets of such dependencies.

Definition 5.5 (Dispersed Graph) Let $g \leq k$ be positive integers. A DAG G is called (g, k) -dispersed if there exists a topological ordering of its nodes such the following holds. Let $[k]$ denote the final k nodes in the ordering of G and let $L_j = [jg, (j+1)g - 1]$ be the j^{th} subinterval. Then $\forall j \in \lceil [k/g] \rceil$ interval L_j has a (g, g) -dependency.

More generally, let $\epsilon \in (0, 1]$. If each interval L_j only has an $(\epsilon g, g)$ -dependency then G is called (ϵ, g, k) -dispersed.

We show that many graphs in the literature consist of a *stack* of dispersed graphs. Our lowerbound on the CC of a dispersed graph grows in the height of this stack. The next definition precisely captures such stacks.

Definition 5.6 (Stacked Dispersed Graphs) A DAG $G = (V, E)$ is called $(\lambda, \epsilon, g, k)$ -dispersed if there exist $\lambda \in \mathbb{N}^+$ disjoint subsets of nodes $\{L_i \subseteq V\}$, each of size k with following two properties.

1. For each L_i there is a path running through all nodes of L_i .
2. Fix any topological ordering of G . For each $i \in [\lambda]$ let G_i be the subgraph of G containing all nodes of G up to the last node of L_i . Then G_i is an (ϵ, g, k) -dispersed graph. We denote the set of $(\lambda, \epsilon, g, k)$ -dispersed graphs by $\mathbb{D}_{\epsilon, g}^{\lambda, k}$.

We are now ready to state and prove the lowerbound on the CC of stacks of dispersed graphs.

Theorem 5.7

$$G \in \mathbb{D}_{\epsilon, g}^{\lambda, k} \Rightarrow \Pi_{cc}^{\parallel}(G) \geq \epsilon \lambda g \left(\frac{k}{2} - g \right).$$

Intuitively we sum the CC of pebbling the last k nodes L_i of each subgraph G_i . For this we consider any adjacent intervals A of $2g$ nodes in L_i . Let p be a path in the $(\epsilon g, g)$ -dependency of the second half of A . Either at least one pebble is always kept on p while pebbling the first half of A (which takes time at least g since a path runs through L_i) or p must be fully pebbled in order to finish pebbling interval A (which also takes time at least g). Either way pebbling A requires an additional CC of g per path in the $(\epsilon g, g)$ -dependency of the second half of A . Since there are $k/2g$ such interval pairs each with ϵg incoming paths in their dependencies we get a total cost for that layer of $kg\epsilon/2$. So the cost for all layer of G is at least $\lambda kg\epsilon/2$. The details (for the more general case when g doesn't divide n) are in Appendix C.

Now that we have our lowerbound for stacks of dispersed graphs it remains to analyze for which parameters various graphs in the literature confirm to this notion. The results are summarized in the theorem below.

Theorem 5.8 [iMHF Constructions Based on Dispersed Graphs]

- If $\lambda, n \in \mathbb{N}^+$ such that $k = n/(\lambda + 1)$ is a power of 2 then it holds that

$$\text{BRG}_{\lambda}^n \in \mathbb{G}_{n, 2} \quad \text{BRG}_{\lambda}^n \in \mathbb{D}_{1, \lceil \sqrt{k} \rceil}^{\lambda, k} \quad \Pi_{cc}^{\parallel}(\text{BRG}_{\lambda}^n) = \Omega\left(\frac{n^{1.5}}{\sqrt{\lambda}}\right)$$

- If $\lambda, n \in \mathbb{N}^+$ such that $n = \bar{n}(\lambda(2c - 1) + 1)$ where $\bar{n} = 2^c$ for some $c \in \mathbb{N}^+$ then it holds that

$$\text{DBG}_{\lambda}^n \in \mathbb{G}_{n, 3} \quad \text{DBG}_{\lambda}^n \in \mathbb{D}_{1, \lceil \sqrt{\bar{n}} \rceil}^{\lambda, \bar{n}} \quad \Pi_{cc}^{\parallel}(\text{DBG}_{\lambda}^n) = \Omega\left(\frac{n^{1.5}}{c\sqrt{c\lambda}}\right)$$

- If $\sigma, \tau \in \mathbb{N}^+$ such that $n = \sigma * \tau$ then with high probability it holds that

$$\text{Lin}_{\tau}^{\sigma} \in \mathbb{G}_{n, 21} \quad \text{Lin}_{\tau}^{\sigma} \in \mathbb{D}_{0.25, \sqrt{\sigma}/2}^{\tau-1, \sigma} \quad \Pi_{cc}^{\parallel}(\text{Lin}_{\tau}^{\sigma}) = \Omega\left(\frac{n^{1.5}}{\sqrt{\tau}}\right).$$

The theorem is proven in three subsections of Appendix B, one per graph.

6 New Memory-Efficient Evaluation Algorithm and Applications

In this section we introduce a new class of memory-efficient evaluation algorithms for non-depth robust DAGs and instantiate attacks on several candidate iMHFs. We show that their asymptotic memory-hardness in the PROM is significantly lower than both the honest algorithm, and even then that shown in [AB16a].

While we instantiate the new algorithm for the case of random DAGs (improving on the hitherto best results of [AB16a]) we believe it very likely to also lead to improved asymptotic attacks on a variety of other iMHFs constructions in the literature. Indeed, the inspiration of the new attack is the algorithm of [AB16a] which has since been shown to result in strong asymptotic attacks on many other iMHF candidates in the literature [AGK⁺16]. Moreover memory complexity of [AB16a] has been shown to be significantly better in actual experiments than its asymptotics might indicate [AB16b].

Review of [AB16a]. The general pebbling attack of Alwen and Blocki [AB16a] made use of a node set $S \subset V$ of size $|S| = e$ such that removing S reduces the depth of the DAG $\text{depth}(G - S) \leq d$. Intuitively keeping pebbles on S compresses the graph in the sense that it can quickly be completely pebbled within d (parallel) steps as any unpebbled path has length $\leq d$. The attack never removes pebbles from nodes in S and its goal is to always pebble node i at time i so as to finish in $n = \text{size}(G)$ steps. To ensure that parents of node i are all pebbled at time i the attack sorts nodes in topological order and them up into chunks of g nodes. Each new chunk is pebbled by a “light phase”. As an invariant, in the first step j at the beginning of a light phase the all parents p of the nodes $[j, j + g - 1]$ with $p < j$ are pebbled. Thus a light phase simply places one pebble on each node in $[j, j + g - 1]$ in sequence over the next g steps.

To guarantee the invariant for a light phase beginning at step g is satisfied the attack makes also runs a “balloon phase” in parallel beginning at step $j - d - 1$ which and runs it for d steps ending just in time for the light phase. Intuitively, during a balloon phase the graph is quickly “decompressed” by greedily re-pebbling everything using the pebbles on nodes in S . Then, in the final step of the balloon phase, all pebbles are removed from the graph except those on nodes in S and those needed for the invariant of the upcoming light phase.

The cost of their attack \mathcal{A} was $\Pi_{cc}^{\parallel}(\mathcal{A}) \leq en + \delta gn + \frac{dn^2}{g}$. Intuitively, the $en + \delta gn$ terms upper bound the pebbling costs of light phases, and the third term upper bounds the cost of all balloon phases — each balloon phase costs at most dn and they need at most $\frac{n}{g}$ balloon phases. Setting $g = \sqrt{nd/\delta}$ to balance the gn and $\frac{dn^2}{g}$ terms they obtained $\Pi_{cc}^{\parallel}(\mathcal{A}) \leq O(en + n^{1.5}\sqrt{d/\delta})^6$.

Recursive Attack. In order to derive the new pebbling strategy we first make some new observations about [AB16a].

1. A balloon phase running during steps $[j - d - 1, j - 1]$ is actually able to pebble *any* subset of nodes in $[1, j - 1]$ (not just those required by the invariant for the next light phases). This can be done with no difference in the balloon phase’s energy complexity from that sketched above as the only difference is in what pebbles are removed in the final step.
2. The starting condition and goal of the [AB16a] algorithm is actually more general than described in that work. It can begin at any node i conditioned on all nodes in $[1, i - 1] \cap S$ already containing a pebbled. In [AB16a] only the special case where $i = 1$ is considered. Moreover, given the above insight about balloon phases, the [AB16a] algorithm can actually be used to pebble any subset of nodes, not just node n as considered in [AB16a].

Armed with these observations, the key insight behind our improved attack is that we may not actually need to pay cost dn for each of the balloon phases. If we could find a second set $S' \supset S$ of $e' > e$ nodes such that $\text{depth}(G - S') \leq d' < d$ then we could potentially reduce the cost of each balloon phase by applying the [AB16a] attack recursively to pebble the parents needed for the next light phase. In general, given a

⁶For some DAGs (e.g., BRG _{λ}) Alwen and Blocki [AB16a] were able to obtain slightly better bounds on $\Pi_{cc}^{\parallel}(\mathcal{A})$ by exploiting instance-specific structure of the particular DAG in their analysis.

function $f(d)$ we can improve the attack for f -reducible DAGs — a DAG that is not $(f(d), d)$ -depth robust for any value of d .

Outline. In the remainder of this section we define f -reducibility. Next we state the main theorem which upper bounds the cumulative complexity of the new attack when applied to any f -reducible graph. We then show that several iMHF candidates are based on f -reducible graphs which gives us a corollary upper-bounding their asymptotic cumulative memory complexity in the PROM. Finally we prove the main theorem (while the missing proofs can be found in the appendix).

Definition 6.1 Let $G = (V, E)$ be a DAG with n nodes and let $f : \mathbb{N} \rightarrow \mathbb{N}$ be a function. We say that G is f -reducible if for every positive integer $n \geq d > 0$ there exists a set $S \subseteq V$ of $|S| = f(d)$ nodes such that $\text{depth}(G - S) \leq d$.

We show that any $f(d) = \tilde{O}(\frac{n}{\sqrt{d}})$ -reducible DAG on n nodes has $\Pi_{cc}^{\parallel}(G) \leq O(n^{1.708})$. In particular, this shows that $\Pi_{cc}^{\parallel}(G) \leq O(n^{1.708})$ for a (n, δ, n) -random DAGs like Argon2i-A. We also show that $\Pi_{cc}^{\parallel}(G) \leq O(n^{1.6181})$ for a $f(d) = \tilde{O}(\frac{n}{d})$ -reducible DAG G , which implies that $\Pi_{cc}^{\parallel}(\text{BRG}_{\lambda}^n) = O(n^{1.6181}) = \Pi_{cc}^{\parallel}(\text{DBG}_{\lambda}^n)$ whenever $\lambda = O(\text{polylog}(n))$. For comparison, the non-recursive attack of Alwen and Blocki [AB16a] demonstrated that $\Pi_{cc}^{\parallel}(G) \leq \tilde{O}(n^{1.75})$ for (n, δ, n) -random DAGs and that $\Pi_{cc}^{\parallel}(G) \leq O(n^{5/3})$ for λ -layered DAGs.

Theorem 6.2 Let G be a f -reducible DAG on n nodes then if $f(d) = \tilde{O}(\frac{n}{d^b})$ for some constant $0 < b \leq 1$ then for any constant $\epsilon > 0$

$$\Pi_{cc}^{\parallel}(G) \leq O(n^{1+a+\epsilon}), \text{ where } a = \frac{1 - 2b + \sqrt{1 + 4b^2}}{2}.$$

Lemma 6.3 states that (n, δ, n) -random DAGs and λ -layered DAGs are f -reducible. The proof of Lemma 6.3, which repeats arguments from Alwen and Blocki [AB16a] for different values of d , can be found in the appendix.

Lemma 6.3 Let $f_b(d) = \tilde{O}(\frac{n}{d^b})$ then

1. Let $\delta = O(\text{polylog}(n))$ then a (n, δ, n) -random DAG is $f_{0.5}$ -reducible with high probability.
2. The Catena DAGs BRG_{λ}^n and DBG_{λ}^n are both f_1 -reducible for $\lambda = O(\text{polylog}(n))$.
3. The Balloon Hashing Linear (and the Double Buffer) graph $\text{Lin}_{\tau}^{\sigma}$ is f_1 -reducible for $\tau = O(\text{polylog}(n))$.

Corollary 6.4 Let $\epsilon > 0$ be any constant

1. Let $\delta = O(\text{polylog}(n))$ then an (n, δ, n) -random DAG G has $\Pi_{cc}^{\parallel}(G) = O(n^{1+\sqrt{1/2}+\epsilon}) \approx O(n^{1.707+\epsilon})$.
2. Both $\Pi_{cc}^{\parallel}(\text{BRG}_{\lambda}^n)$ and $\Pi_{cc}^{\parallel}(\text{DBG}_{\lambda}^n)$ are in $O(n^{1+\frac{\sqrt{5}-1}{2}+\epsilon}) \approx O(n^{1.618+\epsilon})$.
3. $\Pi_{cc}^{\parallel}(\text{Lin}_{\tau}^{\sigma}) = O(n^{1+\frac{\sqrt{5}-1}{2}+\epsilon}) \approx O(n^{1.618+\epsilon})$, where $\text{Lin}_{\tau}^{\sigma}$ has $n = \tau * \sigma$ nodes.

We define a recursive algorithm `RecursiveGenPeb`, which takes as input a DAG G of depth $d_0 = \text{depth}(G)$, a target set $T \subset V(G)$ of nodes to pebble and sets S_1, \dots, S_k of sizes $e_1 = f(d_1) < e_2 = f(d_2) \dots < e_k = f(d_k)$ such that $\text{depth}(G - S_i) \geq d_i$ for each $i \leq k$. `RecursiveGenPeb` returns a pebbling P_1, \dots, P_{2d_0} of G in at most $2d_0$ steps.

We let $G_0 = G$ and $G_{i+1} = G_i - S_{i+1}$. We partition the nodes of G_i according to their depth and further split up sets which are larger than n/d_i . This gives us a partition $D_1^i, \dots, D_{2d_i}^i$ of $V(G_i)$ which satisfies the following properties:

SOURCES: $\text{parents}(D_1^i) = \emptyset$,

TOPOLOGICALLY ORDERED: $\forall j \leq 2d_i - 1 \quad \text{parents}(D_{j+1}^i) \subseteq \bigcup_{y \leq j} D_y^i$, and

MAXIMUM SIZE: $\forall j \leq 2d_i \quad |D_j^i| \leq \frac{n}{d_i}$.

At top level of recursion we closely follow the attack of Alwen and Blocki [AB16a]. Our goal is to pebble G_0 with the target set $T_0 = \text{sinks}(G_0)$ in at most $2d_0$ steps. We ensure that our pebbling maintains the invariant that for pebbling round j we keep pebbles on the set $P_j \supset D_j^0 \cup \text{parents}(\{j+1, \dots, j+g_0\}) \cup (S_1 \cap \bigcup_{y \leq j} D_y^0)$. Thus, during pebbling step j RecursiveGenPeb will pebble all of the nodes in the set D_j^0 .

Light Phases. To save space, during a light phase round i we discard most other pebbles except pebbles on nodes in the set $S_1 \cup D_i^0 \cup \text{parents}(\{i+1, \dots, i+2e_1\})$. We only keep pebbles in the set S_1 and on parents of the next $2e_1$ nodes that we want to pebble. The total cost of all light phases is $2d_0(2\delta + 1)e_1$.

Balloon Phases and Recursion. At some point in time we may find that $D_{i+2d_1+1}^0 \not\subseteq P_i \cup (\bigcup_{y \leq i+2d_1} D_y^0)$ — in $2d_1 + 1$ steps we will want to pebble a node $v \in D_{i+2d_1+1}^0$ but we have already discarded pebble(s) from some of the parent(s) of this node. Thus, to maintain our invariant we will quickly recover the discarded pebbles for all of parents of the next $2e_1$ nodes. Alwen and Blocki [AB16a] used a greedy approach (for d_1 steps pebble everything possible) to accomplish this task (total cost: $d_1 n$). Instead, we call our algorithm RecursiveGenPeb recursively to find a pebbling $(P_1^1, \dots, P_{2d_1}^1)$ of the graph $G_1 = G - S_1$ with the target set $T_1 = \text{parents}(\{i+1, \dots, i+2e_1\})$ in at most $2d_1$ pebbling steps. The recursive pebbling $P_1^1, \dots, P_{2d_1}^1$ will also be divided into balloon phases and light phases, and we will maintain the same invariant until we finish pebbling T_1 (i.e., never discard pebbles from S_2 , always keep pebbles on the parents of the next e_2 nodes we want to pebble, pebble the set $P_i^1 \supset D_i^1$ in round). Balloon phases involve yet another recursive call to RecursiveGenPeb with the graph G_2 and target set $T_2 = \{\text{parents of the next } 2e_2 \text{ nodes we want to pebble in } G_2\}$. At the final level of recursion we use the greedy pebbling strategy (pebble every possible node for $d_k = \text{depth}(G_k)$ steps) which costs at most $d_k n$.

Analysis. We select $n \geq d_0 > d_1 > \dots > d_k$ such that for each $i > 0$ we can find a set S_i of size $|S_i| = e_i = f(d_i) = \tilde{O}(n^{a_i})$ such that $\text{depth}(G - S_i) \leq d_i$. Here, the sequence $\{a_i\}_{i=1}^\infty$ is defined as follows

$$a_1 = a = \frac{1 - 2b + \sqrt{1 + 4b^2}}{2}, \quad \text{and} \quad a_{i+1} = 1 + \frac{(a-1)(1-a_i)}{b}.$$

We obtain the following recurrence relationship from our attack

$$\Pi_{cc}^{\parallel}(G_i) \leq (2\delta + 1)e_{i+1}2d_i + \frac{n}{e_{i+1}}\Pi_{cc}^{\parallel}(G_{i+1}).$$

In particular, our pebbling of G_i takes at most $2d_i$ steps and we keep pebbles on at most $(2\delta + 1)e_{i+1}$ nodes during light phases. We need to execute at most $\frac{n}{e_{i+1}}$ balloon phases and each balloon phase costs at most $\Pi_{cc}^{\parallel}(G_{i+1})$ to complete. For a sufficiently large constant $k > k_\epsilon$ we have $d_k < n^{\epsilon/2}$ so that $\Pi_{cc}^{\parallel}(G_k) \leq O(n^{1+\epsilon/2})$. By exploiting several key properties of the sequence $\{a_i\}_{i=1}^\infty$ we can unroll the recurrence to show that

$$\Pi_{cc}^{\parallel}(G_0) \leq O(n^{1+a+\epsilon}).$$

7 Open Questions

We conclude with several open questions for future research.

- We showed that for some constant $c \geq 0$ we can find a DAG G on n nodes with $\Pi_{cc}^{\parallel}(G) \geq cn^2 / \log(n)$ and $\text{indeg}(G) = 2$. While this result is asymptotically optimal the constant terms are relevant for practical applications to iMHFs. How big can this constant c be? Can we find explicit constructions of constant-indegree, $(c_1 n / \log(n), c_2 n)$ -depth robust DAGs that match these bounds?

- Provide tighter upper and lower bounds on $\Pi_{cc}^{\parallel}(G)$ for Argon2i-B [BDKJ16], the most recent version of Argon2i which was submitted to IRTF for standardization.
- Another interesting direction concerns understanding the cumulative pebbling complexity of generic graphs. Given a graph G is it computationally tractable to (approximately) compute $\Pi_{cc}^{\parallel}(G)$? An efficient approximation algorithm for $\Pi_{cc}^{\parallel}(G)$ would allow us to quickly analyze candidate iMHF constructions. Conversely, as many existing iMHF constructions are based on fixed random graphs, [BDK16, CGBS16] showing that approximating such a graphs complexity is hard would provide evidence that an adversary will likely not be able to leverage properties of the concrete instance to improve their evaluation strategy for the iMHF. Indeed, it may turn out that the most effective way to construct depth-robust graphs with good constants is via a randomized construction.
- Another open challenge is to develop a data-dependent MHF f with parallel cost $\Pi_{cc}^{\parallel}(f) = \Omega(n^2)$. Alwen and Blocki [AB16a] show that this goal is impossible for iMHFs, but data-dependent MHFs could potentially satisfy both properties. However, at this time it is not known whether any dMHF candidate (e.g., `script` [Per09, PJ12]) satisfy this property.

References

- [AB16a] Joël Alwen and Jeremiah Blocki. Efficiently computing data-independent memory-hard functions. In *Advances in Cryptology CRYPTO'16*. Springer, 2016.
- [AB16b] Joël Alwen and Jeremiah Blocki. Towards practical attacks on argon2i and balloon hashing. Cryptology ePrint Archive, Report 2016/759, 2016. <http://eprint.iacr.org/2016/759>.
- [ABW03] Martín Abadi, Michael Burrows, and Ted Wobber. Moderately hard, memory-bound functions. In *Proceedings of the Network and Distributed System Security Symposium, NDSS 2003, San Diego, California, USA*, 2003.
- [ACK⁺16a] Joël Alwen, Binyi Chen, Chethan Kamath, Vladimir Kolmogorov, Krzysztof Pietrzak, and Stefano Tessaro. On the complexity of `script` and proofs of space in the parallel random oracle model. Cryptology ePrint Archive, Report 2016/100, 2016. <http://eprint.iacr.org/>.
- [ACK⁺16b] Joël Alwen, Binyi Chen, Chethan Kamath, Vladimir Kolmogorov, Krzysztof Pietrzak, and Stefano Tessaro. On the complexity of `script` and proofs of space in the parallel random oracle model. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 358–387. Springer, Heidelberg, May 2016.
- [ACK⁺16c] Joël Alwen, Binyi Chen, Chethan Kamath, Vladimir Kolmogorov, Krzysztof Pietrzak, and Stefano Tessaro. On the complexity of `Script` and proofs of space in the parallel random oracle model. Cryptology ePrint Archive, Report 2016/100, 2016. <http://eprint.iacr.org/2016/100>.
- [ACP⁺16] Jol Alwen, Binyi Chen, Krzysztof Pietrzak, Leonid Reyzin, and Stefano Tessaro. `script` is maximally memory-hard in the parallel random oracle model. Manuscript, 2016.
- [AGK⁺16] Jol Alwen, Peter Gai, Chethan Kamath, Karen Klein, Georg Osang, Krzysztof Pietrzak, Leonid Reyzin, Michal Rolnek, and Michal Rybr. On the memory-hardness of data-independent password-hashing functions. Cryptology ePrint Archive, Report 2016/783, 2016. <http://eprint.iacr.org/2016/783>.
- [AS15] Joël Alwen and Vladimir Serbinenko. High Parallel Complexity Graphs and Memory-Hard Functions. In *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing, STOC '15*, 2015. <http://eprint.iacr.org/2014/238>.

- [BDK15] Alex Biryukov, Daniel Dinu, and Dmitry Khovratovich. Fast and tradeoff-resilient memory-hard functions for cryptocurrencies and password hashing. Cryptology ePrint Archive, Report 2015/430, 2015. <http://eprint.iacr.org/2015/430>.
- [BDK16] Alex Biryukov, Daniel Dinu, and Dmitry Khovratovich. Argon2 password hash. Version 1.3, 2016. <https://www.cryptolux.org/images/0/0d/Argon2.pdf>.
- [BDKJ16] Alex Biryukov, Daniel Dinu, Dmitry Khovratovich, and Simon Josefsson. The memory-hard Argon2 password hash and proof-of-work function. Internet-Draft draft-irtf-cfrg-argon2-00, Internet Engineering Task Force, March 2016.
- [BK15] Alex Biryukov and Dmitry Khovratovich. Tradeoff cryptanalysis of memory-hard functions. Cryptology ePrint Archive, Report 2015/227, 2015. <http://eprint.iacr.org/>.
- [CGBS16] Henry Corrigan-Gibbs, Dan Boneh, and Stuart Schechter. Balloon hashing: Provably space-hard hash functions with data-independent access patterns. Cryptology ePrint Archive, Report 2016/027, Version: 20160601:225540, 2016. <http://eprint.iacr.org/>.
- [Cha73] Ashok K. Chandra. Efficient compilation of linear recursive programs. In *SWAT (FOCS)*, pages 16–25. IEEE Computer Society, 1973.
- [Coo73] Stephen A. Cook. An observation on time-storage trade off. In *Proceedings of the Fifth Annual ACM Symposium on Theory of Computing, STOC '73*, pages 29–33, New York, NY, USA, 1973. ACM.
- [DFKP15] Stefan Dziembowski, Sebastian Faust, Vladimir Kolmogorov, and Krzysztof Pietrzak. Proofs of space. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 585–605. Springer, Heidelberg, August 2015.
- [DGN03] Cynthia Dwork, Andrew Goldberg, and Moni Naor. On memory-bound functions for fighting spam. In *Advances in Cryptology - CRYPTO 2003, 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003, Proceedings*, volume 2729 of *Lecture Notes in Computer Science*, pages 426–444. Springer, 2003.
- [DKW11a] Stefan Dziembowski, Tomasz Kazana, and Daniel Wichs. Key-evolution schemes resilient to space-bounded leakage. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 335–353. Springer, Heidelberg, August 2011.
- [DKW11b] Stefan Dziembowski, Tomasz Kazana, and Daniel Wichs. One-time computable self-erasing functions. In Yuval Ishai, editor, *TCC*, volume 6597 of *Lecture Notes in Computer Science*, pages 125–143. Springer, 2011.
- [DNW05] Cynthia Dwork, Moni Naor, and Hoeteck Wee. Pebbling and proofs of work. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 37–54. Springer, Heidelberg, August 2005.
- [EGS75] Paul Erdős, Ronald L. Graham, and Endre Szemerédi. On sparse graphs with dense long paths. Technical report, Stanford, CA, USA, 1975.
- [FLW13] Christian Forler, Stefan Lucks, and Jakob Wenzel. Catena: A memory-consuming password scrambler. *IACR Cryptology ePrint Archive*, 2013:525, 2013.
- [HJO⁺16] Brett Hemenway, Zahra Jafargholi, Rafail Ostrovsky, Alessandra Scafuro, and Daniel Wichs. Adaptively secure garbled circuits from one-way functions. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 149–178. Springer, Heidelberg, August 2016.

- [HP70] Carl E. Hewitt and Michael S. Paterson. Record of the project mac conference on concurrent systems and parallel computation. chapter Comparative Schematology, pages 119–127. ACM, New York, NY, USA, 1970.
- [JW16] Zahra Jafargholi and Daniel Wichs. Adaptive security of yao’s garbled circuits. Cryptology ePrint Archive, Report 2016/814, 2016. <http://eprint.iacr.org/2016/814>.
- [Kal00] Burt Kaliski. Pkcs# 5: Password-based cryptography specification version 2.0. 2000.
- [LT82] Thomas Lengauer and Robert E. Tarjan. Asymptotically tight bounds on time-space trade-offs in a pebble game. *J. ACM*, 29(4):1087–1130, October 1982.
- [MMV13] Mohammad Mahmoody, Tal Moran, and Salil P. Vadhan. Publicly verifiable proofs of sequential work. In Robert D. Kleinberg, editor, *ITCS 2013*, pages 373–388. ACM, January 2013.
- [NB⁺15] Arvind Narayanan, Joseph Bonneau, , Edward W Felten, Andrew Miller, and Steven Goldfeder. *Bitcoin and Cryptocurrency Technology (manuscript)*. 2015. Retrieved 8/6/2015.
- [Per09] C. Percival. Stronger key derivation via sequential memory-hard functions. In *BSDCan 2009*, 2009.
- [PHC] Password hashing competition. <https://password-hashing.net/>.
- [PJ12] Colin Percival and Simon Josefsson. The scrypt password-based key derivation function. 2012.
- [PR80] Wolfgang J. Paul and Rüdiger Reischuk. On alternation II. A graph theoretic approach to determinism versus nondeterminism. *Acta Inf.*, 14:391–403, 1980.
- [RD16] Ling Ren and Srinivas Devadas. Proof of space from stacked bipartite graphs. Cryptology ePrint Archive, Report 2016/333, 2016. <http://eprint.iacr.org/>.
- [Sch82] Georg Schnitger. A family of graphs with expensive depth reduction. *Theor. Comput. Sci.*, 18:89–93, 1982.
- [Sch83] Georg Schnitger. On depth-reduction and grates. In *24th Annual Symposium on Foundations of Computer Science, Tucson, Arizona, USA, 7-9 November 1983*, pages 323–328. IEEE Computer Society, 1983.
- [SS78] John E. Savage and Sowmitri Swamy. Space-time trade-offs on the fft algorithm. *IEEE Transactions on Information Theory*, 24(5):563–568, 1978.
- [SS79a] John E. Savage and Sowmitri Swamy. Space-time tradeoffs for oblivious interger multiplications. In Hermann A. Maurer, editor, *ICALP*, volume 71 of *Lecture Notes in Computer Science*, pages 498–504. Springer, 1979.
- [SS79b] Sowmitri Swamy and John E. Savage. Space-time tradeoffs for linear recursion. In Alfred V. Aho, Stephen N. Zilles, and Barry K. Rosen, editors, *POPL*, pages 135–142. ACM Press, 1979.
- [Tom78] Martin Tompa. Time-space tradeoffs for computing functions, using connectivity properties of their circuits. In *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing*, STOC ’78, pages 196–204, New York, NY, USA, 1978. ACM.
- [Val77] Leslie G. Valiant. Graph-theoretic arguments in low-level complexity. In Jozef Gruska, editor, *Mathematical Foundations of Computer Science 1977, 6th Symposium, Tatranska Lomnica, Czechoslovakia, September 5-9, 1977, Proceedings*, volume 53 of *Lecture Notes in Computer Science*, pages 162–176. Springer, 1977.

A Memory-Hard Functions

We define MHFs in the Parallel Random Oracle Model (pROM) of [AS15]. For this we first define the model and associated complexity notions and then fix the exact notion of MHF.

The Parallel Random Oracle Model. We consider an arbitrary repeatedly invoked algorithm \mathcal{A} executing in the parallel random oracle model (pROM) [AS15] of computation where we make states between invocations explicit as follows. At invocation $i \in \{1, 2, \dots\}$ algorithm \mathcal{A} is given the state (bit-string) σ_{i-1} it produced at the end of the previous invocation. Next \mathcal{A} can make a batch of calls $\mathbf{q}_i = (q_{1,i}, q_{2,i}, \dots)$ to the *fixed input length* random oracle H (a.k.a. an ideal compression function). Then it receives the response from H and can perform arbitrary computation before finally outputting an updated state σ_i . The initial state σ_0 contains the input to the computation which terminates once a special final state is produced by \mathcal{A} . Apart from the explicit states σ the algorithm may keep no other state between invocations. For an input x and coins r we denote by $\mathcal{A}(x; r; H)$ the corresponding (deterministic) execution of \mathcal{A} . We say that \mathcal{A} is *sequential* if in no execution does it ever make a batch of queries \mathbf{q} with $|\mathbf{q}| > 1$.

The *cumulative memory complexity* (CMC) is defined to be

$$\text{cmc}(\mathcal{A}) = \mathbb{E}_H \left[\max_{x,r} \sum_i |\sigma_i| \right]$$

where $|\sigma|$ is the bit-length of state σ , the expectation is taken over the choice of H and $\max_{x,r}$ denotes the maximum over all inputs and coins of \mathcal{A} .

We also need the following two worst-case complexity notions. The *time complexity* (TC) $\text{time}(\mathcal{A})$ is the maximum running time of \mathcal{A} in any execution (over all choices of x , r and H). Similarly, the *space complexity* (SC) is the largest state it ever outputs in any execution.

$$\text{space}(\mathcal{A}) = \max_{x,r,H} \{ \max_i |\sigma_i| \}.$$

We remark that SC and TC are somewhat stricter than usual since we maximise over all choices of H . However, this can only help us as these measures are used as worst case estimates of the complexity for the honest party and we ask that an MHF has reasonable upper-bounds on their values.

An oracle function f is a function over strings which depends on the choice of H . Let $\mathbb{A}_{f,m,q}$ be the set of pROM algorithms which compute f on $m \in \mathbb{N}^+$ arbitrary (distinct) inputs making at most q queries to H . Then the *amortized cumulative memory complexity* (aCMC) of f is defined to be

$$\text{cmc}_{m,q}(f) = \min \left\{ \frac{\text{cmc}(\mathcal{A})}{n} : n \in [m], \mathcal{A} \in \mathbb{A}_{f,n,q} \right\}.$$

We comment on two differences between the above complexity notions and those in [AS15] (and why they do not prevent us from using the results in that work).

- In the definition of CMC above we maximize over all coins of \mathcal{A} instead of including them in the expectation. This is with out loss of generality for the aCMC of a function since hardcoding the coins which minimize the aCMC of any \mathcal{A} has at least as much CMC as the expected value for random coins.
- The aCMC of [AS15] is also parametrized by the minimum success probability ϵ of any algorithm computing f . Instead, for reasons of exposition, in this work we restrict ourselves to the special case when $\epsilon = 1$.

Memory-Hard Functions. As observed in [AS15] the aCMC of a function provides a good lower-bound on the amortized AT-complexity of that function. Thus the following definition captures the intuition of a memory-hard function in the pROM.

Definition A.1 (Memory-Hard Function) Let $\{f_{\sigma,\tau}\}_{\sigma,\tau \in \mathbb{N}^+}$ be a family of (oracle) functions and \mathcal{N} be a sequential pROM algorithm which, on input (σ, τ, x) , outputs $f_{\sigma,\tau}(x)$ in time at most $\tau\sigma$ using space at most σ . Then $F = (\{f_{\sigma,\tau}\}, \mathcal{N})$ is an (h, g, t) -memory-hard function (for up to m instances and q queries) if it has memory-hardness at least h , memory-gap at most g and throughput at least t (all functions of σ and τ).

$$\text{cmc}_{m,q}(f_{\sigma,\tau}) \geq h(\sigma, \tau) \qquad \frac{\text{space}(\mathcal{N}) * \text{time}(\mathcal{N})}{\text{cmc}_{m,q}(f_{\sigma,\tau})} \leq g(\sigma, \tau) \qquad \frac{\text{space}(\mathcal{N})}{\text{time}(\mathcal{N})} \geq t(\sigma, \tau).$$

In practice, we may content ourselves with families of infinite size but which do not contain a member for possible $k \in \mathbb{N}^+$ as long as the family is not too sparse (e.g. we have a function for all powers of 2).

Using the following theorem from [AS15] the results in this work imply MHF's with various desirable properties.

Theorem A.2 ([AS15]) Let $G \in \mathbb{G}_{n,\delta}$, P be a sequential pebbling of G with $(\Pi_t(P), \Pi_s(P)) = (z_n, s_n - 1)$ and let H be a random oracle with $w > 13$ bits of output. Fix any $m, q \in \mathbb{N}^+$ subject to the following (reasonable) pair of constraints.

- Not too many copies of f are computed: $m \leq 2^{w-2}/n$.
- Not too many oracle queries are made: $q \leq 2^{w/2}$.

Then there exists an oracle function f and pROM evaluation algorithm \mathcal{N} for f with

$$\text{time}(\mathcal{N}) = z_n \qquad \text{space}(\mathcal{N}) = w * s_n \qquad \text{cmc}_{m,q}(f) \geq \frac{w * \Pi_{cc}^{\parallel}(G)}{4}.$$

In particular this theorem shows that in order to construct an MHF with both memory-cost σ and time-cost τ parameters it suffices to construct a family of DAGs for every size (with high CC) together with matching sequential pebbblings. For simplicity we state the theorem for DAGs with a single source and sink but it can be easily extended to the more general case.⁷

Corollary A.3 (High CC Graphs Imply MHFs) Let $\{G_n \in \mathbb{G}_{n,\delta_n}\}_{n=1}^{\infty}$ be a family of DAGs each with a single source and sink and let H be a random oracle with $w > 13$ bits of output. For each n let P_n be a sequential pebbling of G_n where $(\Pi_t(P_n), \Pi_s(P_n)) = (z_n, s_n - 1)$. Then there exists a (h, g, t) -MHF with

$$h(\sigma, \tau) \geq \frac{w * \tau * \Pi_{cc}^{\parallel}(G_{\sigma})}{4} \qquad g(\sigma, \tau) \leq \frac{4z_{\sigma} * s_{\sigma}}{\Pi_{cc}^{\parallel}(G_{\sigma})} \qquad t(\sigma, \tau) \geq \frac{w * s_{\sigma}}{z_{\sigma} \tau}$$

PROOF. The idea is simple. Fix any σ and τ . To obtain oracle function connect τ copies of the DAG G_{σ} in a chain and let $f_{\sigma,\tau}$ be the MHF given by Theorem A.2. The corresponding sequential pebbling $P_{\sigma,\tau}$ is simply the pebbling P_{σ} repeated τ times for each copy of G_{σ} with the following caveat. Whenever a pebble is placed on the source node of any copy of G_{σ} it is only removed once no more children of that node will be pebbled. Thus $\text{space}(P_{\sigma,\tau}) \leq \text{space}(P_{\sigma}) + w = w * s_{\sigma}$ and $\text{time}(P_{\sigma,\tau}) = \tau * \text{time}(P_{\sigma}) = \tau z_{\sigma}$. Finally, it is easy to see that $\Pi_{cc}^{\parallel}(G_{\sigma,\tau}) = \tau * \Pi_{cc}^{\parallel}(G_{\sigma})$ since any pebbling of $G_{\sigma,\tau}$ with CC less than $\tau * \Pi_{cc}^{\parallel}(G_{\sigma})$ would have to pebble at least one copy of G_{σ} with less than $\Pi_{cc}^{\parallel}(G_{\sigma})$ which is a contradiction. \square

Two remarks are in order.

- In practice one would probably want a stronger connection between copies of G_{σ} . For example one could connect the last σ nodes pebbled by P_{σ} in one copy of G_{σ} to the first σ nodes pebbled by P_{σ} in the next copy of G_{σ} . This would affect neither the time nor space complexity of the honest evaluation algorithm but would potentially increase the concrete (though not asymptotic) memory-hardness of $f_{\sigma,\tau}$ which would also result in a smaller memory-gap. For the purpose of this work though the simple construction in the proof suffices as it has the same asymptotic behaviour.

⁷Moreover any DAG can easily be extended to be of this form with no penalty to the CC as a function of its size.

- Estimating the effect of requiring a pebble on the source of each internal copy of G_σ to potentially require the space complexity to grow by 1 is extremely pessimistic. To the best of our knowledge the \mathcal{N} algorithm for all MHF constructions in the literature as well the sequential pebbling of all DAGs in this work already keep such a pebble there (rather than re-pebble the source repeatedly). Thus the space complexity of the sequential pebbling when composing those DAGs would not grow by 1. Never the less, rather than make the corollary seem less general than it is (by requiring such a property from P) we have opted for its current form. In particular the difference is both asymptotically, and practically speaking, imaterial.

B Dispersed Graphs

B.1 Catena Bit Reversal

We begin with the Catena Bit Reversal graph. We first briefly recall the properties of the BRG_λ^n construction, relevant to our proof, summarized in the following lemma which follows easily from the definition of BRG_λ^n in [FLW13, Def. 8 & 9].

For this we describe the “bit-reversal” function from which the graph derives its name. Let $k \in \mathbb{N}^+$ such that $c = \log_2 k$ is an integer. On input $x \in [k]$ the *bit-reversal function* $\mathbf{br}(\cdot) : [k] \rightarrow [k]$ returns y such that the binary representation of x using c bits is the reverse of the binary representation of y using c bits.

Lemma B.1 (Catena Bit Reversal Graph) *Let $\lambda, n \in \mathbb{N}^+$ be such that $k = n/(\lambda + 1)$ is a power of 2. Let $G = \text{BRG}_\lambda^n$ be the Catena Bit Reversal graph. Then the following holds:*

1. G has n nodes.
2. Number them in topological order with the set $[0, n - 1]$ and $\forall i \in [\lambda]$ let node set $L_i = [ik, (i + 1)k - 1]$. A path runs through all nodes in each set L_i .
3. Node $ki + x \in L_i$ has an incoming edge from $k(i - 1) + \mathbf{br}(x) \in L_{i-1}$.

We are ready to state and prove the lowerbound.

Lemma B.2 *It holds that $\text{BRG}_\lambda^n \in \mathbb{D}_{1, \sqrt{k}}^{\lambda, k}$ where $k = \frac{n}{(\lambda + 1)}$ and $\Pi_{cc}^\parallel(\text{BRG}_\lambda^n) = \Omega\left(\frac{n^{1.5}}{\sqrt{\lambda}}\right)$.*

Proof of Lemma B.2. Let $G = \text{BRG}_\lambda^n$ and set $k = n/(\lambda + 1)$, $c = \log_2 k$ and $g = \sqrt{k}$. By construction c is an integer. For simplicity assume c is even and so $g \in \mathbb{N}^+$.⁸ Number the nodes of G according to (the unique) topological order with the set $[0, n - 1]$. It suffices to show that for all $i \in [\lambda]$ the subgraph G_i consisting of nodes $[(i + 1)k - 1]$ is (g, k) -dispersed (with probability $\epsilon = 1$). If this holds then Theorem 5.7 immediately implies that $\Pi_{cc}^\parallel(\text{BRG}_\lambda^n) = \Omega\left(\frac{n^{1.5}}{\sqrt{\lambda}}\right)$.

Fix any $i \in [\lambda]$ and let the nodes of layer L_i be the nodes $[ik, (i + 1)k - 1]$. We represent the nodes of L_i with bit strings of length c . For $j \in [k/2g]$ let $L_{i,j} = \{ik + 2gj + 1, \dots, ik + 2gj + 2g\}$ be the j 'th interval of length $2g$ in layer L_i and let $R = \{ik + 2gj + g + 1, \dots, ik + 2gj + 2g\}$ be the second half of $L_{i,j}$. We will show that there are g node-distinct paths terminating in R with all other nodes in layer L_{i-1} . Define set $S_x = [s + y - (g - 2), s + y]$ where $y = \mathbf{br}(x)$ and $s = (i - 1)2^{(c+1)}$. The next three properties are follow immediately from Lemma B.1 and they imply the lemma.

- $\forall x \in R$ it holds that $S_x \subset L_{i-1}$.
- $\forall x \in R$ there is a path of length g going through the nodes of S_x and ending in x .
- \forall distinct $x, x' \in R$ sets S_x and $S_{x'}$ are disjoint.

⁸The odd case is identical but with messy but inconsequential rounding terms.

B.2 Catena Double Butterfly

Next we focus on the Catena Double Butterfly graph. We begin with a lemma (which follows immediately by inspection of the Catena Double Butterfly definition [FLW13, Def. 10 & 11]) summarizing the properties of that construction relevant to our proof.

Lemma B.3 (Catena Double Butterfly Graph) *Let $\lambda, n \in \mathbb{N}^+$ such that $n = \bar{n}(\lambda(2c - 1) + 1)$ where $\bar{n} = 2^c$ for some $c \in \mathbb{N}^+$. Then the Catena Double Butterfly Graph DBG_λ^n consists of a stack of λ subgraphs such that the following holds.*

1. *The graph DBG_λ^n has n nodes in total.*
2. *The graph DBG_λ^n built as a stack of λ subgraphs $\{G_i\}_{i \in [\lambda]}$ each of which is a superconcentrator. In the unique topological ordering of DBG_λ^n denote the first and final \bar{n} nodes of each G_i as $L_{i,0}$ and $L_{i,1}$ respectively. Then there is a path running through all nodes in each $L_{i,1}$.*
3. *Moreover, for any $i \in [\lambda]$ and subsets $S \subset L_{i,0}$ and $T \subset L_{i,1}$ with $|S| = |T| = h \leq \bar{n}$ there exist h node disjoint paths p_1, \dots, p_h of length $2c$ from S to T .*

We are ready to state and prove the lowerbound.

Lemma B.4 *Let $\lambda, n \in \mathbb{N}^+$ such that $n = \bar{n}(\lambda(2c - 1) + 1)$ with $\bar{n} = 2^c$ for some $c \in \mathbb{N}^+$. It holds that $\text{DBG}_\lambda^n \in \mathbb{D}_{1,g}^{\lambda,n}$ for $g = \lceil \sqrt{\bar{n}} \rceil$ and $\Pi_{cc}^\parallel(\text{DBG}_\lambda^n) = O\left(\frac{n^{1.5}}{c\sqrt{c\lambda}}\right)$.*

Proof of Lemma B.4. Let $G = \text{DBG}_\lambda^n$ and let $G_1, G_2, \dots, G_\lambda$ be the subgraphs of G described in Lemma B.3. We will show that each G_i is (g, \bar{n}) -dispersed for $g = \lceil \sqrt{\bar{n}} \rceil$. Fix arbitrary $i \in [\lambda]$ and L_1 be the last \bar{n} nodes in the (the unique) topological ordering of G_i . We identify the nodes in L_1 with the set $\{1\} \times [\bar{n}]$ such that the second component follows their topological ordering. Let $\bar{g} = \lfloor \bar{n}/g \rfloor$ and for each $j \in [\bar{g}]$ let $L_{1,j} = \{(1, ig + x) : x \in [0, g - 1]\}$ ik+2gj+2g We will show that $L_{1,j}$ has a (g, g) -dependency.

Let L_0 be the first \bar{n} nodes of G_i which identify with the set $\{0\} \times [\bar{n}]$ (again with the second component respecting their ordering). Notice that for $n > 1$ and $g = \lceil \sqrt{\bar{n}} \rceil$ it holds that $g(g - 2c + 1) \leq n$. Thus the set $S = \{(0, i(g - 2c + 1)) : i \in [g]\}$ is fully contained in L_0 . Property (3) of Lemma B.3 implies there exist g node disjoint paths from S to $L_{1,j}$ of length $2c$. In particular $L_{1,j}$ has a $(2c, g)$ -dependency.

We extend this to a (g, g) -dependency. Let path p , beginning at node $\langle 0, v \rangle \in S$, be a path in the $(2c, g)$ -dependency of $L_{1,j}$. Prepend to p the path traversing

$$\langle (0, v - (g - 2c - 1)), \langle 0, v - (g - 2c - 2) \rangle, \dots, \langle 0, v \rangle$$

to obtain a new path p^+ of length g . As this is a subinterval of L_0 property (2) of Lemma B.3 implies this prefix path always exists. Moreover since any paths $p \neq q$ in a $(g, 2c)$ -dependency of $L_{1,i}$ are node disjoint they must, in particular, also begin at distinct nodes $\langle 0, v_p \rangle \neq \langle 0, v_q \rangle$ in S . But by construction of S any such pair of nodes is separated by $g - 2c$ nodes. In particular paths p^+ and q^+ are also node disjoint and so by extending all paths in a $(2c, g)$ -dependency we obtain a (g, g) -dependency for $L_{1,i}$. This concludes the first part of the lemma.

It remains to lowerbound $\Pi_{cc}^\parallel(\text{DBG}_\lambda^n)$ using Theorem 5.7.

$$\Pi_{cc}^\parallel(\text{DBG}_\lambda^n) \geq \lambda g \left(\frac{k}{2} - g \right) \geq \lambda \lceil \sqrt{\bar{n}} \rceil \left(\frac{\bar{n}}{2} - \lceil \sqrt{\bar{n}} \rceil \right) = \lambda \sqrt{\bar{n}} \left(\frac{\bar{n}}{2} - \sqrt{\bar{n}} \right) - O(\bar{n}) = \Omega(\lambda \bar{n}^{1.5}) = \Omega\left(\frac{n^{1.5}}{c\sqrt{c\lambda}}\right).$$

B.3 Balloon Hashing Linear Graph

The graph $G = \text{Lin}_\tau^\sigma$ is a pseudo-randomly constructed τ -layered graph with $\text{indeg}(G) = 21$. It is defined as follows:

$G = (V, E)$ has n nodes $V = \{0, \dots, n-1\}$, and G contains a path $0, 1, 2, \dots, n-1$ running through V .

For $0 \leq i < \tau$ let $L_i = [i\sigma, (i+1)\sigma - 1]$ denote the i 'th layer. For each node $x \in L_i$, with $i > 0$, we select 20 nodes $y_1, \dots, y_n \in L_{i-1}$ (uniformly at random) and add the directed edges $(y_1, x), \dots, (y_{20}, x)$ to E .

Lemma B.5 *If $\sigma, \tau \in \mathbb{N}^+$ such that $n = \sigma * \tau$ then with high probability it holds that*

$$\text{Lin}_\tau^\sigma \in \mathbb{G}_{n,21} \quad \text{Lin}_\tau^\sigma \in \mathbb{D}_{0.25, \sqrt{\sigma}/2}^{\tau, \sigma} \quad \Pi_{\text{cc}}^\parallel(\text{Lin}_\tau^\sigma) = \Omega\left(\frac{n^{1.5}}{\sqrt{\tau}}\right).$$

PROOF. (sketch) It suffices to show that $\text{Lin}_\tau^\sigma \in \mathbb{D}_{0.25, \sqrt{\sigma}/2}^{\tau-1, \sigma}$. By Theorem 5.7 it immediately follows that

$$\Pi_{\text{cc}}^\parallel(\text{Lin}_\tau^\sigma) \geq \frac{(\tau-1)\sqrt{\sigma}/2}{4} (\sigma/2 - \sqrt{\sigma}/2) = \Omega\left(\frac{n^{1.5}}{\sqrt{\tau}}\right).$$

Consider layer L_i ($\tau \geq i > 0$) and given let $S_x = \{x, \dots, x + \sqrt{\sigma}/2 - 1\} \subset L_i$ denote a segment of $\sqrt{\sigma}/2$ consecutive nodes in L_i . Without loss of generality we suppose that each node in S_x only has one randomly chosen parent in L_{i-1} — adding additional edges can only improve dispersity. We partition L_{i-1} into $\sqrt{\sigma}$ intervals of length \sqrt{s} . We say that an interval $\{u, \dots, u + \sqrt{s} - 1\} \subset L_{i-1}$ is covered by S_x if there is a directed edge (y, v) from $y \in \{u + \sqrt{s}/2, \dots, u + \sqrt{s} - 1\}$ (the last half of the interval) to some $v \in S_x$. In this case the path $u, u+1, \dots, y, v$ has length $\geq \sqrt{s}/2$ and this path will not intersect the corresponding paths from any of the other (disjoint) intervals in L_{i-1} (recall that we are assuming that $v \in S_x$ only has one parent in L_{i-1}). The probability that an interval $\{u, \dots, u + \sqrt{s} - 1\} \subset L_{i-1}$ is covered by S_x is at least

$$1 - \left(1 - \frac{\sqrt{\sigma}/2}{\sigma}\right)^{\sqrt{\sigma}} \approx 1 - \sqrt{1/e}.$$

Thus, in expectation we will have at least $\mu = \sqrt{\sigma} \left(1 - \sqrt{1/e}\right) \geq 0.39 \times \sqrt{\sigma}$ node disjoint paths of length $\sqrt{\sigma}/2$ ending in S_x . Standard concentration bounds imply that we will have at least $\sqrt{\sigma}/4$ such paths with high probability. □

C Missing Proofs

Reminder of Theorem 3.2. *Let DAG $G = (V, E)$ be (e, d) -depth-robust and let $S, T \subset V$ such that*

$$|S| \leq \sigma n \quad , \quad |\text{ancestors}_{G-S}(T)| \geq \tau n \quad \text{and} \quad T \cap S = \emptyset$$

Then the cost of pebbling $G - S$ with target set T is $\Pi_{\text{cc}}^\parallel(G - S, T) > (e - \sigma n)(d - (1 - \tau)n)$.

Proof of Theorem 3.2. Set $\bar{e} = (e - \sigma n)$ and $\bar{d} = d - (1 - \tau)n$. As G is (e, d) -depth-robust the DAG $G - S$ is (\bar{e}, \bar{d}) -depth-robust: assume for the sake of contradiction there exists $B \subset V(G - S)$ of size $|B| \leq \bar{e}$ s.t. $G - (S \cup B)$ has no path of length d , but then $S' = S \cup B$ is a set of size $|S'| = |S| + |B| \leq \sigma n + \bar{e} = e$ where $G - S'$ has no path of length d contradicting G 's (e, d) -depth robustness.

Now consider the subgraph G_T of $G - S$ which contains only the vertices $\text{ancestors}_{G-S}(T)$, note that $\Pi_{\text{cc}}^\parallel(G - S, T) = \Pi_{\text{cc}}^\parallel(G_T)$. Below we'll show that G_T is (\bar{e}, \bar{d}) -depth-robust, using this and Corollary 3.1 we further get

$$\Pi_{\text{cc}}^\parallel(G - S, T) = \Pi_{\text{cc}}^\parallel(G_T) > \bar{e} \cdot \bar{d} = (e - \sigma n)(d - (1 - \tau)n)$$

which proves the theorem.

We still must show that G_T is (\bar{e}, \bar{d}) -depth-robust. Assume for contradiction that there exists a set $A \subset V(G_T)$, $|A| \leq \bar{e}$ s.t. $G_T - A$ has no path of length \bar{d} . As G_T is a subgraph of $G - S$ and $G - S$ has $(1 - \tau)n$ vertices more than G_T we have no path of length $d = \bar{d} + (1 - \tau)n$ in $G - (S \cup A)$. This contradicts the (\bar{e}, d) -depth-robustness of $G - S$.

Reminder of Corollary 3.3. For some constants $c_1, c_2 > 0$ there exists an infinite family of DAGs $\{G_{n,\delta} \in \mathbb{G}_{n,\delta}\}_{n=1}^\infty$ with $\delta \leq c_1 \log(n)$ and $\Pi_{cc}^\parallel(G) \geq c_2 n^2$.

This is optimal in the sense that for any family $\{\delta_n \in [n]\}_{n=1}^\infty$ and $\{J_n \in \mathbb{G}_{n,\delta_n}\}_{n=1}^\infty$ it holds that $\Pi_{cc}^\parallel(J_n) \in O(n^2)$. Moreover if $\delta_n = o(\log(n)/\log \log(n))$ then $\Pi_{cc}^\parallel(J_n) = o(n^2) = o(\Pi_{cc}^\parallel(G_n))$.

Proof of Corollary 3.3. Take $\{G_n \in \mathbb{G}_{n,c_3 \log(n)}\}_{n=1}^\infty$ to be the family of DAGs from Theorem 2.2. Now the first and second statements follow immediately from Theorem 3.2 and the observation that $\Pi_{cc}^\parallel(J_n) \leq n^2$ for any n node DAG J_n . The final statement is based on the simple observation that $\Pi_{cc}^\parallel(J_n) = o(n^2)$ whenever $\delta_n = o(\log(n)/\log \log(n))$ because any such DAG is $(o(n), o(n/\delta_n))$ -reducible. We can see this by applying Lemma C.1, due to Valiant [Val77], $3 \log(\delta_n)$ times to obtain a set S of at most

$$e = |S| \leq \frac{3 \log(\delta_n) n \delta_n}{\log(n) - 2 \log(\delta_n)} = o(n)$$

nodes such that $d = \text{depth}(G) \leq 2^{-3 \log(\delta_n)} n = o(n/\delta_n^2)$. Now by Theorem 2.5 we have $\Pi_{cc}^\parallel(G) = O(ne + n\sqrt{dn\delta_n}) = o(n^2)$. \square

Lemma C.1 ([Val77] Extension) Given a DAG G with m edges and depth $\text{depth}(G) \leq d = 2^i$ there is a set of m/i edges s.t. by deleting them we obtain a graph of depth at most $d/2$.

Reminder of Lemma 5.3. For any $e \geq \sqrt{n}$ any any $\delta \geq 2$ a (n, δ, n) -random DAG will be $(e, \Omega(\frac{n^2}{e^2 \log(n)}), \frac{n}{100e})$ -block depth robust with probability $1 - o(n^{-3})$.

Proof of Lemma 5.3. Let G be a random (n, δ, n) -random DAG $G = (V, E)$ with nodes $V = [n]$. Fix an arbitrary integer $m \in [n]$ and set $n' = \lfloor n/m \rfloor$. We will define a DAG G_m called the meta-graph of G . For this we use the following sets. For all $i \in [n']$ let $S_i = [(i-1)m + 1, im] \subseteq V$. Moreover we denote the first and last thirds respectively of S_i with

$$S_i^F = [(i-1)m + 1, (i-1)m + \lfloor m/3 \rfloor] \subseteq S_i \quad S_i^L = [(i-1)m + \lceil 2m/3 \rceil + 1, im].$$

We define the meta-graph $G_m = (V_m, E_m)$ as follows:

Nodes: V_m contains one node v_i per set S_i . We call v_i the *simple node* and S_i its *meta-node*.

Edges: If the end of a meta-node is connected to the beginning of another meta-node we connect their simple nodes.

$$V_m = \{v_i : i \in [n']\} \quad E_m = \{(v_i, v_j) : E \cap (S_i^L \times S_j^F) \neq \emptyset\}.$$

Lemma 5.3 now follows from the next two claims.

Claim C.2 If G_m is (e, d) -depth robust then G is $(e/(1 + \lceil \frac{n}{100em} \rceil), dm/3, \frac{n}{100e})$ -block depth robust.

PROOF. (sketch) Notice meta-nodes are disjoint subsets. Fix any set $S \subseteq V$ of size e . The set

$$N\left(S, \frac{n}{100e}\right) = \bigcup_{v \in S} \left\{v - \frac{n}{100e} + 1, \dots, v\right\}$$

can intersect at most $e(1 + \lceil \frac{n}{100em} \rceil)$ meta-nodes because for each $v \in S$ the set $\{v - \frac{n}{100e} + 1, \dots, v\}$ intersects at most $(1 + \lceil \frac{n}{100em} \rceil)$ meta-nodes. Remove the simple nodes corresponding to $N(S, \frac{n}{100e})$ from

G_m and there remains a path ϕ' in G_m of length d . Thus, after removing S from G there must remain a path ϕ going through the middle thirds of the d meta-nodes corresponding to ϕ' . In particular we get that

$$\text{length}(\phi) \geq \frac{\text{length}(\phi') * m}{3} \geq \frac{dm}{3}.$$

□

Claim C.3 *With high probability G_m is $(n'/10, m/(200 \log(n)))$ -depth robust.*

PROOF. (sketch) Divide G_m into $d = m/(100 \log(n))$ layers L_1, \dots, L_d each containing n'/d meta-nodes. We first observe that (with high probability) each fixed pair of layers L_i and L_j is connected with a γ -bipartite expander graph: Fix any pair of layers L_i and L_j ($i < j$) and let $S \subset L_i$ and $V \subset L_j$ be such that $|S| > \gamma n'/d$ and $|V| > \gamma n'/d$ then $|\text{parents}(V) \cap S| > 0$ — for a small value γ (say $\gamma \leq 0.1$). To see why this expansion property holds (whp) we note that any two meta-nodes are connected with probability at least $\frac{m^2}{3^2 n}$. Thus, a meta-node in L_i is connected to $\frac{n'}{d} \times \frac{m^2}{9n} \geq 10 \log(n)$ random meta-nodes in L_j . Thus, except with very small probability, L_i and L_j will be connected with a γ -bipartite expander graph and we can union bound over all $\binom{d}{2}$ pairs $i < j$ to show that this holds for every pair with high probability⁹.

Suppose that the above expansion property holds, and suppose that we remove a set S of $n'/10$ meta-nodes from G_m . Call a layer L_i good if $|S \cap L_i| \leq \frac{n'}{5d}$. By Markov's inequality we have at least $d/2$ good layers. Let $Y_1, \dots, Y_{d/2}$ denote $d/2$ good layers. Now we claim that there is a path traversing through every good layer (hence, $\text{depth}(G_m - S) \geq d/2$). Let $R_1 = Y_1 - S$ denote the set of meta-nodes in the first good layer which remain in $G_m - S$. In general, $R_{i+1} = \text{parents}(Y_{i+1} - S) \cap (R_i - S)$ denotes the set of meta-nodes in good layer Y_{i+1} that are reachable by a path in $G_m - S$. Now a simple inductive argument shows that $|R_i| \geq \frac{4n'}{5d}$ for all $i \leq d/2$.

□

From Claim C.2 and Claim C.3 it follows that G is $\left(\frac{\lfloor n/m \rfloor}{10(1 + \lceil \frac{n}{100em} \rceil)}, m^2/(600 \log(n)), \frac{n}{100e}\right)$ -depth robust.

Setting $m = 20n/e$ we obtain the desired result. □

Reminder of Theorem 5.2. *Let G be a (n, δ, n) -random DAG then, except with probability $o(n^{-3})$, we have*

$$\Pi_{cc}^{\parallel}(G) = \tilde{\Omega}\left(n^{5/3}\right).$$

Proof of Theorem 5.2. Let G be a (n, δ, n) -random DAG, and let $G_1 = G - \{n/2 + 1, \dots, n\}$ denote the subgraph on the first $n/2$ nodes. We observe that G_1 is itself a $(n/2, \delta, n/2)$ -random DAG.

By setting $e = n^{2/3}$ Lemma 5.3 we can find some fixed constant $c > 0$ such that our graph G_1 is $(n^{2/3}, cn^{2/3}/\log(n), n^{1/3}/100)$ -block depth robust except with probability $o(n^{-3})$.

Now consider any fixed segment $B_v = \{v, v + 1, \dots, v + 10\tau n^{2/3} \log(n) - 1\} \subset \{n/2 + 1, \dots, n\}$. That is, B_v is a segment of $10\tau n^{2/3} \log(n)$ consecutive nodes in the last half of G . Let $B_v^{first} = \{v, v + 1, \dots, v + 5n^{2/3}\tau \log(n) - 1\}$ and $B_v^{last} = \{v + 5n^{2/3}\tau \log(n), \dots, v + 10n^{2/3}\tau \log(n) - 1\}$. We first note that we can select $\tau = O(1)$ such that for every segment $B_u = \{u, u + 1, \dots, u + n^{2/3}\} \subset V(G_1)$ of $n^{2/3}/100$ consecutive nodes in G_1 there is an edge from B_u to B_v^{last} with high probability.

Now we consider the cost incurred while pebbling B_v . Let t_1 (resp. t_2, t_3) denote the first time at which we place a pebble on node v (resp. node $v + 5n^{2/3}\tau \log(n) - 1$, node $v + 10n^{2/3}\tau \log(n) - 1$).

Claim C.4 *Let P_0, \dots, P_t denote a pebbling of G then*

$$\sum_{i \in [t_1, t_3]} |P_i| = \tilde{\Omega}\left(n^{4/3}\right).$$

⁹It is well known that a random bipartite graph with degree $d = \Omega(1)$ is an expander with good probability (e.g., see Erdős et al. [EGS75]). In our case we have degree $d = \Omega(\log(n))$ so the bipartite graph will be a γ expander with overwhelming probability.

PROOF. There are two cases:

1. For every $i \in [t_1, t_2]$ we have $|P_i| \geq e$ pebbles on G . Total Cost: $\sum_{i \in [t_1, t_2]} |P_i| \geq e(t_2+1-t_1) = \Omega(n^{4/3})$.
2. For some $t' \in [t_1, t_2]$ we have at most $|P_{t'}| < e$ pebbles on G .

In the second case we define S_0, \dots, S_{d-1} where

$$S_i = \bigcup_{j \in [t', t_3]: j \equiv i \pmod{d}} (P_j \cap V(G_1)) .$$

We can find some i such that $|S_i| \leq \sum_{i \in [t_1, t_3]} |P_i|/d$. Now we observe that, because (whp) there is a backedge from B_v^{last} to every set B_u of $n^{2/3}/100$ consecutive nodes in G_1 , every node in

$$V(G_1) - \bigcup_{v \in S_i} \left\{ v - \frac{n^{1/3}}{100} + 1, \dots, v \right\}$$

must be pebbled at some point in time in the interval $[t', t_3]$. This means that

$$\text{depth} \left(G - \bigcup_{v \in S_i} \left\{ v - \frac{n^{1/3}}{100} + 1, \dots, v \right\} \right) \leq d .$$

By the $(n^{2/3}, cn^{2/3}/\log(n), n^{1/3}/100)$ -block depth robustness of G_1 it follows that $\sum_{i \in [t', t_3]} |P_i| = \tilde{\Omega}(n^{4/3})$. \square

It follows from Claim C.4 that the total cost of a pebbling is at least $\frac{n}{2|B_v|} \tilde{\Omega}(n^{4/3}) = \tilde{\Omega}(n^{5/3})$. \square

Reminder of Theorem 5.7.

$$G \in \mathbb{D}_{\epsilon, g}^{\lambda, k} \Rightarrow \Pi_{cc}^{\parallel}(G) \geq \epsilon \lambda g \left(\frac{k}{2} - g \right) .$$

Proof of Theorem 5.7. We number the nodes of $G = (V, E)$ in a topological order with the set $[n]$. Let $\{L_i \subseteq V\}$ be the λ disjunct node subsets described in Definition 5.6. For any L_i and $j \in [\lfloor k/2g \rfloor]$ let $L_{i,j}$ be the j^{th} interval of $2g$ consecutive nodes in L_i . That is if $L_i = [a, a+k-1]$ then $L_{i,j} = [a+i2g, a+(i+1)2g-1]$. Let G_i denote the subgraph of G consisting of the nodes up to the end of L_i . That is G consists exactly of the node set $[a + (i+1)2g - 1]$ and edges of G between those nodes.

Let $P = (P_1, P_2, \dots)$ be a legal pebbling of G and let \bar{t}_v denote the first time step at which node v is pebbled by P . We use the following shorthand:

$$c_{i,j} := \sum_{i=\bar{t}_\alpha}^{\bar{t}_z} |P_i| \qquad c_i := \sum_{i=\bar{t}_\alpha}^{\bar{t}_\omega} |P_i|$$

where a and z are the first and last nodes of $L_{i,j}$ and α and ω are the first and last nodes of L_i respectively. In particular $\text{cc}(P) \geq \sum_{i \in [\lambda]} c_j$ and our goal is to lowerbound this sum. The theorem follows from the next claim.

Claim C.5 *For all $i \in [\lambda]$ and $j \in [\lfloor k/2g \rfloor]$ we have $c_{i,j} \geq \epsilon g^2$.*

In particular, from Claim C.5, it immediately follows that

$$c_i \geq \left\lfloor \frac{k}{2g} \right\rfloor (\epsilon g^2) \geq \epsilon g \left(\frac{k}{2} - g \right) .$$

Therefore it follows that

$$\Pi_{cc}^{\parallel}(G) \geq \sum_{i \in [\lambda]} c_i \geq \epsilon \lambda g \left(\frac{k}{2} - g \right).$$

It remains only to prove Claim C.5. \square

Proof of Claim C.5. Fix any i and j as in the claim. Let L and R be the first and second halves of $L_{i,j}$ and let a and y be the first and last nodes of L respectively and let z denote the last node of R . As G_i is (ϵ, g) -dispersed there exists a set Γ of ϵg node disjoint paths each of length g , terminating R and with no other nodes in L_i . Let $p \in \Gamma$ be such a path.

Now either p contains (at least) one pebble during all time steps in $[\bar{t}_a, \bar{t}_y]$ or not. If so p contributes at least g to $c_{i,j}$ with pebbles not lying on any of other $p' \in \Gamma$. If not then in order to pebble z , the last node of R , all of p must also be pebbled between steps $[\bar{t}_a, \bar{t}_z]$. Thus again p contributes at least g to $c_{i,j}$ with pebbles not lying on any other $p' \in \Gamma$.

Summing over all paths in Γ we get that $c_{i,j} \geq \epsilon g^2$ which concludes the proof of the claim and theorem. \square

Reminder of Lemma 6.3. Let $f_b(d) = \tilde{O}\left(\frac{n}{d^b}\right)$ then

1. Let $\delta = O(\text{polylog}(n))$ then a (n, δ, n) -random DAG is $f_{0.5}$ -reducible with high probability.
2. The Catena DAGs BRG_{λ}^n and DBG_{λ}^n are both f_1 -reducible for $\lambda = O(\text{polylog}(n))$.
3. The Balloon Hashing Linear (and the Double Buffer) graph $\text{Lin}_{\tau}^{\sigma}$ is f_1 -reducible for $\tau = O(\text{polylog}(n))$.

The proof of Lemma 6.3 closely follows arguments from Alwen and Blocki [AB16a], who proved these DAGs were (e, d) -reducible for specific values of e and d . Because the attack of Alwen and Blocki [AB16a] was non-recursive they only focused on proving (e, d) -reducible for the values e, d which optimized the quality of their attack.

Proof of Lemma 6.3. (sketch) We first consider an arbitrary $\lambda = O(\text{polylog}(n))$ -layered DAG G . This includes Catena DAGs BRG_{λ}^n and DBG_{λ}^n as well as Balloon Hashing Linear (and the Double Buffer) graph $\text{Lin}_{\tau}^{\sigma}$ (with $\tau = \lambda = O(\text{polylog}(n))$). Let d be given and let $e = (\lambda + 1)n/d$. For simplicity assume that $e, n/(\lambda + 1)$ and $d/(\lambda + 1)$ are integers¹⁰. Let $S = \{i \times d/(\lambda + 1) : i \leq g\}$ and observe that S has size $|S| = e = \tilde{O}(n/d)$. Thus, to show that G is f_1 reducible it suffices to show that $\text{depth}(G - S) \leq d$. Define $L_i = \{i \times n/(\lambda + 1) + 1, \dots, (i + 1) \times n/(\lambda + 1)\}$ for $0 \leq i \leq \lambda$. We note that any path in $G - S$ can remain on layer L_i for at most $d/(\lambda + 1)$ steps before moving to a higher layer and there are $\lambda + 1$ layers. Thus, the maximum length of any path is $(\lambda + 1)d/(\lambda + 1) = d$.

Next we consider with a (n, δ, n) -random DAG G on nodes $\{1, \dots, n\}$. Let d be given and let $g = n/\sqrt{d}$. For simplicity assume that \sqrt{d} and g are integers¹¹. Let $S_1 = \{i \times \sqrt{d} : i \leq g\}$ and observe that S_1 has size $|S_1| = g$. Define $L_i = \{i \times g + 1, \dots, i \times g + g\}$. We call a node $v \in L_i$ good if $\text{parents}(v) \cap L_i = \emptyset$ and we let $B_i = \{v \in L_i : \text{parents}(v) \cap L_i \neq \emptyset\}$ denote the set of all bad vertices in L_i . Finally, we let $S = S_1 \cup \bigcup_{i=0}^{\sqrt{d}-1} B_i$. It is easy to verify that $\text{depth}(G - S) \leq d$ because any path in $G - S$ can remain on layer L_i for at most \sqrt{d} steps before moving up to a higher layer and there are \sqrt{d} layers. Thus, the maximum length of any path is $\sqrt{d}^2 = d$. It is easy to see that $\mathbb{E}[|B_i|] \leq \frac{|B_i|}{i+1} = \frac{\delta g}{i+1}$ — let x_v denote the indicator random variable for the event $v \in B_i$ then $\Pr[x_v = 1] \leq \frac{1}{i+1}$. Thus, $\mathbb{E}[|S|] \leq g + \delta g \sum_{i=1}^g i^{-1} = O\left(\frac{\delta n \log n}{\sqrt{d}}\right)$. Furthermore, standard concentration bounds imply that $|S| - 2\mathbb{E}[|S|] \leq 0$ except with very small probability. Thus, a (n, δ, n) -random DAG G is $f_{0.5}$ -reducible with high probability. \square

Reminder of Theorem 6.2. Let G be a f -reducible DAG on n nodes then if $f(d) = \tilde{O}\left(\frac{n}{d^b}\right)$ for some constant $0 < b \leq 1$ then for any constant $\epsilon > 0$

$$\Pi_{cc}^{\parallel}(G) \leq O\left(n^{1+a+\epsilon}\right), \text{ where } a = \frac{1 - 2b + \sqrt{1 + 4b^2}}{2}.$$

¹⁰This allows us to simply presentation by ignoring insignificant rounding terms.

¹¹This allows us to simply presentation by ignoring insignificant rounding terms.

Proof of Theorem 6.2. Let c be a constant such that $f(d) \leq \frac{cn \log^c n}{d^b}$. We define a sequence $\{a_i\}_{i=1}^\infty$ as follows

$$a_1 = a = \frac{1 - 2b + \sqrt{1 + 4b^2}}{2}, \quad \text{and} \quad a_{i+1} = 1 + \frac{(a-1)(1-a_i)}{b}.$$

We remark that a is the (positive) solution to the equation $x^2 + (2b-1)x - b = 0$. When we run our recursive attack `RecursiveGenPeb` we will set $d_i = \tilde{O}\left(n^{\frac{1-a_i}{b}}\right)$ such that $e_i = f(d_i) = n^{a_i} = g_{i-1}$. The following claims about the sequence $\{a_i\}_{i=1}^\infty$ will also be useful when analyzing the cost of our recursive attack.

Claim C.6 For $i \geq 1$ we have

$$0 \leq a_i \leq 1, \quad a_{i+1} \geq a_i, \quad \text{and} \quad a_{i+1} + \frac{1-a_i}{b} = 1 - a_{i+1} + a_{i+2} + \frac{1-a_{i+1}}{b}.$$

Claim C.7

$$\lim_{i \rightarrow \infty} a_i = 1, \quad \text{and} \quad \lim_{i \rightarrow \infty} \sum_{j=1}^i (1-a_j) = a.$$

By Claim C.7 for any $\epsilon > 0$ we can find a value k_ϵ such that $1 - a_k \leq \epsilon$ and for all $k > k_\epsilon$.

We run `RecursiveGenPeb`, see Algorithm 1, setting $T = \{n\}$, $k = k_{\epsilon/2} + 1$ and defining $\bar{S} = S_1, \dots, S_k$, $\bar{g} = g_0, g_1, \dots, g_k$ and $\bar{d} = d_0 = n, d_1, \dots, d_k$ as follows: for each $i \leq k$ we let S_i denote a set of size $|S_i| = e_i \leq n^{a_i}$ such that $\text{depth}(G - S_i) \leq d_i = f^{-1}(e_i) \leq c \left(\frac{n}{n^{a_i}}\right)^{1/b} \log^c n$. We also set $g_{i-1} = e_i = n^{a_i}$.

Let $T(k)$ denote the cost of a balloon phase at the final level k of the recursion. In this case a balloon phase lasts for at most $2d_k$ steps so the total cost is at most $T(k) \leq 2d_k n \leq 2cn \left(\frac{n}{n^{a_k}}\right)^{1/b} \log^c n \leq 2cn^{1+\epsilon/2} \log^c n$. In general, we have the following recurrence

$$\begin{aligned} T(i) &\leq e_{i+1}d_i + 2\delta g_i d_i + \frac{n}{g_i} T(i+1) = (2\delta + 1)d_i e_{i+1} + \frac{n}{e_{i+1}} T(i+1) \\ &\leq (2\delta + 1)cn^{a_{i+1} + \frac{1-a_i}{b}} \log^c n + n^{1-a_{i+1}} T(i+1). \end{aligned}$$

In particular, a level- i balloon phase lasts for $2d_i$ rounds and in every step we keep $e_{i+1} + 2\delta g_i$ pebbles on the graph — corresponding to the nodes in S_{i+1} and the parents of the next $2g_i$ nodes that we want to pebble. We make at most $\frac{n}{g_i}$ recursive calls to level- $i+1$ balloon phases — each of cost $\leq T(i+1)$.

Unrolling the recurrence, and applying Claim C.6, we get that

$$\begin{aligned} T(i) &\leq (2\delta + 1)cn^{a_{i+1} + \frac{1-a_i}{b}} \log^c n + n^{1-a_{i+1}} \left((2\delta + 1)cn^{a_{i+2} + \frac{1-a_{i+1}}{b}} \log^c n + n^{1-a_{i+2}} T(i+2) \right) \\ &\leq 2(2\delta + 1)cn^{1 + \frac{a(1-a_i)}{b}} \log^c n + n^{2-a_{i+1}-a_{i+2}} (T(i+2)). \end{aligned}$$

In particular,

$$\begin{aligned}
\Pi_{cc}^{\parallel}(\mathcal{A}) &\leq e_1 n + 2\delta g_0 n + \frac{n}{g_0} T(1) \\
&\leq (2\delta + 1)n^{1+a_1} \log^c n + n^{1-a} \left((2\delta + 1)n^{a_2 + \frac{1-a}{b}} \log^c n \right) + \frac{n^2}{g_0 g_1} T(2) \\
&= 2(2\delta + 1)n^{1+a_1} \log^c n + \frac{n^2}{g_0 g_1} T(2) \\
&\leq (k+1)(2\delta + 1)cn^{1+\frac{a(1-a)}{b}} \log^c n + T(k) \prod_{j=0}^k \frac{n}{g_j} \\
&\leq (k+1)(2\delta + 1)cn^{1+a} \log^c n + T(k)n^{k+1-\sum_{i=0}^k a_{i+1}} \\
&\leq (k+1)(2\delta + 1)cn^{1+a} \log^c n + n^{1+\epsilon/2} n^{\sum_{i=0}^k (1-a_{i+1})} \\
&\leq (k+1)(2\delta + 1)cn^{1+a} \log^c n + n^{1+\epsilon/2+\sum_{i=1}^{\infty} (1-a_i)} \\
&\leq k(\delta + 1)cn^{1+a+\epsilon/2} \log^c n \\
&= O(n^{1+a+\epsilon}) .
\end{aligned}$$

□

Reminder of Claim C.6. For $i \geq 1$ we have

$$0 \leq a_i \leq 1, \quad a_{i+1} \geq a_i, \quad \text{and} \quad a_{i+1} + \frac{1-a_i}{b} = 1 - a_{i+1} + a_{i+2} + \frac{1-a_{i+1}}{b} .$$

Proof of Claim C.6. We first two statements by induction. Clearly, $0 \leq a = a_1 \leq 1$. Now for $i \geq 1$ we have

$$a_{i+1} = 1 + \frac{(a-1)(1-a_i)}{b} = 1 - \left(\frac{(1-a)}{b} \right) (1-a_i) \leq 1 ,$$

because b , $(1-a)$ and $(1-a_i)$ are all positive values. Similarly, we also have

$$a_{i+1} - a_i = 1 - \left(\frac{(1-a)}{b} \right) (1-a_i) - a_i = (1-a_i) \left(1 + \frac{(1-a)}{b} \right) \geq 0 .$$

Thus, $a_{i+1} \geq a_i \geq 0$.

For the third statement we observe that $a_{i+1} + \frac{1-a_i}{b} = \frac{a \cdot a_{i+1} - 1}{a-1}$ and that $1 - a_{i+1} + a_{i+2} + \frac{1-a_{i+1}}{b} = 2 - a_{i+1} + \frac{a(1-a_{i+1})}{b}$. Now we claim that $2 - a_{i+1} + \frac{a(1-a_{i+1})}{b} - \frac{a \cdot a_{i+1} - 1}{a-1} = 0$. Multiplying both sides by $(a-1)b$ we get $2(a-1)b - a_{i+1}(a-1)b + a(a-1)(1-a_{i+1}) - ab \cdot a_{i+1} + b = 0$. Refactoring we get $-b(1-a_{i+1}) + a^2(1-a_{i+1}) + a((2b-1)(1-a_{i+1})) = 0$. Dividing by $(1-a_{i+1})$ we get $a^2 + (2b-1)a - b = 0$, which is true by definition of a . □

Reminder of Claim C.7.

$$\lim_{i \rightarrow \infty} a_i = 1, \quad \text{and} \quad \lim_{i \rightarrow \infty} \sum_{j=1}^i (1-a_j) = a .$$

Proof of Claim C.7. We first note that $0 < \left(\frac{1-a}{b}\right) < 1$, and that

$$\begin{aligned}
1 - a_{i+1} &= \frac{(1-a)(1-a_i)}{b} \\
&= \left(\frac{1-a}{b}\right) (1-a_i) \\
&= \left(\frac{1-a}{b}\right)^2 (1-a_{i-1}) \\
&= \dots \\
&= \left(\frac{1-a}{b}\right)^i (1-a_1).
\end{aligned}$$

Therefore, $\lim_{i \rightarrow \infty} a_i = 1 - \lim_{i \rightarrow \infty} \left(\frac{1-a}{b}\right)^i (1-a_1) = 1$. Similarly,

$$\begin{aligned}
\sum_{j=1}^i (1-a_j) &= \sum_{j=0}^{i-1} \left(\frac{1-a}{b}\right)^j (1-a_1) \\
&= (1-a) \sum_{j=0}^{i-1} \left(\frac{1-a}{b}\right)^j.
\end{aligned}$$

Thus, $\lim_{i \rightarrow \infty} \sum_{j=1}^i (1-a_j) = \frac{1-a}{1-\frac{1-a}{b}} = \frac{b-ba}{b-1+a} = a$, where the last equality follows because a was chosen so that $a^2 + (2b-1)a - b = 0$. \square

C.1 Tighter Bounds on Memory-Hardness

As an example of the power of our results we can immediately improve on the analysis of the high CC graph of [AS15]. In that work, first a graph $G \in \mathbb{G}_{n,\delta}$ with indegree $\delta \leq \log^3(n)$ is given and it is show to have CC approximately $n^2/\log(n)$. Next the indegree is reduced to obtain G' at a cost of δ^3 to the CC. That is G' has size $N = n * \delta$ and CC of roughly $N^2/\log^{10}(N)$.

In contrast Lemma 4.1 only loses a factor of δ in the CC when reducing the indegree. Moreover, the indegree of the depth-robust graphs of [MMV13] used in the construction of [AS15] actually have indegree $\log^2(n)$ times some polyloglog function of n and so the same is true G . Ignoring polyloglog factors and setting $g = n/\log^3(n)$ in Theorem 2.5 we see that G must be at least (e, d) -depth-robust for some $e = \Omega(n/\log(n))$ and $d = \Omega(n/\log^4(n))$. So we can reduce the indegree of G to 2 using Lemma 4.1 with $\gamma = \delta$ and we obtain a graph $H \in G_{n\delta,2}$ of size $N = 2n\delta$ which is $(e, d\delta)$ -depth-robust. In particular Corollary 3.1 shows that the CC of H is at least $\Omega(n^2/\log^3(n)) = \Omega(N^2/\log^7(N))$.

The above analysis uses the results from [AS15] as a blackbox. In fact, if we look into the construction the graph $G_{n,\delta}$ we will observe that is consists of a stack of $\log \log(n)$ depth-robust graphs from the construction of [MMV13]. Corollary 3.1 directly implies that $\Pi_{cc}^{\parallel}(G_{n,\delta}) \geq \frac{cn^2}{(\log \log(n))^2}$ for some constant $c > 0$ (in fact this is just the cost of pebbling the top layer). Lemma 4.1 allows us to reduce the indegree $G_{n,\delta}$ at a cost of $\delta \leq \log^2(n) \text{polylog} \log(n)$. Thus, we obtain a a constant indegree graph on N nodes with $\Pi_{cc}^{\parallel}(G) \geq \frac{cN^2}{\log^2(n) \text{polylog} \log(n)}$.

Algorithm 1: RecursiveGenPeb ($G, k, \bar{S}, \bar{g}, \bar{d}, T$)

Arguments : $G = (V, E), k, S_1 \subseteq \dots \subseteq S_k \subseteq V, g_0, g_1, \dots, g_k$ s.t. $g_{j-1} \in [2 \cdot \text{depth}(G - S_j), |V|]$,
 d_0, d_1, \dots, d_k s.t. $d_j \geq \text{depth}(G - S_j)$ and $T \subseteq V$

Local Variables: $n = |V|$, a node-partition $\emptyset = D_0, D_1, \dots, D_{2d_0} \subseteq V$ s.t. $|D_i| \leq \frac{n}{d_0}$ and
 $\text{parents}(D_{i+1}) \subseteq \bigcup_{j=0}^i D_j$

Output : $P_1, \dots, P_{2d_0} \subseteq V$

```
1  $d \leftarrow \max\{j : T \cap D_j \neq \emptyset\}$  // Need  $d \leq 2d_0$  steps to pebble nodes in  $T$ .
2  $P_0, \dots, P_{2d_0-d} \leftarrow \emptyset$ 
3 if  $k=0$  then // Greedy Pebble
4   for  $i = 1$  to  $d$  do
5      $P_{2d_0-d+i} \leftarrow D_i \cup P_{2d_0-d+i-1}$ 
6   end
7   Return  $P_1, \dots, P_{2d_0}$ 
8 else
9    $A_1, \dots, A_{2d_0} \leftarrow \emptyset$ 
10  for  $i = 1$  to  $d$  do
11     $j_g \leftarrow \min\{j > i : |\bigcup_{r=i+1}^j D_r| \geq g\}$ 
12     $j_{2g} \leftarrow \min\{j > i : |\bigcup_{r=i+1}^j D_r| \geq 2g\}$ 
13     $N_g \leftarrow \text{parents}\left(\bigcup_{r=i+1}^{j_g} D_r\right) \cap \left(\bigcup_{r=0}^{i-1} D_r\right)$ 
14     $N_{2g} \leftarrow \text{parents}\left(\bigcup_{r=i+1}^{j_{2g}} D_r\right) \cap \left(\bigcup_{r=0}^{i-1} D_r\right)$ 
15     $K \leftarrow S_1 \cup N_{2g} \cup T$  // Do not discard pebbles in  $K$ 
16     $A \leftarrow \bigcup_{j=i}^{i+2d_1} A_j$  // Pebbles already scheduled to be added in next  $2d_1$  rounds
17    if  $N_g \not\subseteq D_i \cup P_{i-1} \cup A$  then // Balloon Phase
18       $\bar{S}' \leftarrow S_2, \dots, S_k$ 
19       $\bar{g}' \leftarrow g_1, \dots, g_k$ 
20       $\bar{d}' \leftarrow d_1, \dots, d_k$ 
21       $T' \leftarrow N_{2g} - P_{i-1} - D_i - A$ 
22       $B_1, \dots, B_{2d_1} \leftarrow \text{RecursiveGenPeb}(G - S_1, k - 1, \bar{S}', \bar{g}', \bar{d}', N_{2g})$  // Add  $B_1, \dots, B_{2d_1}$ 
23      // to pebbling schedule
24      for  $j = 1$  to  $2d_1$  do
25         $A_{i+j-1} \leftarrow A_{i+j-1} \cup B_j$ 
26      end
27       $P_{2d_0-d+i} \leftarrow D_i \cup A_i \cup (P_{2d_0-d+i-1} \cap K)$ 
28    else // Light Phase
29       $P_{2d_0-d+i} \leftarrow D_i \cup A_i \cup (P_{2d_0-d+i-1} \cap K)$ 
30    end
31  end
32  Return  $P_1, \dots, P_{2d_0}$ 
33 end
```
